arXiv:1410.5993v1 [cs.LO] 22 Oct 2014

The Relative Succinctness and Expressiveness of Modal Logics Can Be Arbitrarily Complex

Henning Schnoor

Institut für Informatik, Christian-Albrechts-Universität zu Kiel, 24098 Kiel, Germany
henning.schnoor@email.uni-kiel.de

Abstract. We study the relative succinctness and expressiveness of modal logics, and prove that these relationships can be as complex as any countable partial order. For this, we use two uniform formalisms to define modal operators, and obtain results on succinctness and expressiveness in these two settings. Our proofs are based on formula size games introduced by Adler and Immerman and bisimulations.

Introduction

Modal logics of different flavours play an important role in computer science, especially as specification languages ([HM92,FHMV95,ABvdT10,BdRV01]). Therefore, the study of expressiveness and succinctness of modal and other logics have received much attention: In [GKPS95], the succinctness of different formalisms to define knowledge bases was compared. In [Wil99], it was shown that CTL$^+$ is exponentially more succinct than CTL, i.e., in the translation from CTL* to CTL, an exponential blow-up in formula size cannot be avoided. This result was later strengthened in [AI03]. The techniques of the latter paper, *formula size games*, were applied to modal logics in [FvdHIK13] and [vdHI14]. They show that, in particular, there are modal logics $\mathcal{L}_1$ and $\mathcal{L}_2$ such that $\mathcal{L}_1$ is exponentially more succinct than $\mathcal{L}_2$ and vice versa. (This seemingly contradictory statement says that some properties are more efficiently expressed in $\mathcal{L}_1$, and some in $\mathcal{L}_2$).

This result raises several questions: Are there arbitrary large sets of modal logics, where each logic is exponentially more succinct than all of the others? Are there arbitrarily long sequences of modal logics of strictly increasing succinctness? More generally, can the "succinctness"-relationships between modal logics be arbitrarily complex?

Formally, let $\leq_{poly}$ be the relation between modal logics such that $\mathcal{L}_1 \leq_{poly} \mathcal{L}_2$ if for every $\mathcal{L}_1$-formula, there is an equivalent $\mathcal{L}_2$-formula of polynomial size. The results from the above-mentioned [vdHI14] imply that $\leq_{poly}$ is not a linear order, but clearly, $\leq_{poly}$ is reflexive and transitive. Does $\leq_{poly}$ have any other standard properties in addition to reflexivity and transitivity?

We answer the above questions by showing that $\leq_{poly}$ can be as complex as any countable partial order. More precisely, for any partial order $\leq_S$ on a countable set S, we exhibit a family of modal logics $(\mathcal{L}_s)_{s\in S}$, all equally expressive, such that $\leq_{poly}$ on $(\mathcal{L}_s)_{s\in S}$ behaves exactly like $\leq_S$ in the following sense: If

$s_1 \leq_S s_2$, then $\mathcal{L}_{s_1} \leq_{poly} \mathcal{L}_{s_2}$ and $\mathcal{L}_{s_1}$ is exponentially more succinct than $\mathcal{L}_{s_2}$ otherwise. In particular, there is indeed an infinite set of modal logics where each logic is exponentially more succinct than every other, and there is an infinite sequence of modal logics, each strictly more succinct than the previous one. For the related question of expressiveness, we get analogous results: If $\leq_{expr}$ is defined as $\mathcal{L}_1 \leq_{expr} \mathcal{L}_2$ if for every $\mathcal{L}_1$-formula, there is some equivalent $\mathcal{L}_2$-formula (regardless of the size), then $\leq_{expr}$ can be as complex as any countable partial order in exact same sense as above.

To prove our results, we use a uniform way to define modal logics. We consider two different ways to define generalized (multi-)modal operators:

1. "One-Step" modal operators, similar to the ones defined in [GPT87], only "look one step ahead in the structure." Such an operator $\Box_f$ is given by the Boolean function f that "selects" a successor world w' of w based on the R_i-relationships between w and w' for each accessibility relation R_i. As there are only finitely many Boolean functions of a given arity, this only allows to prove our main result for finite partial orders S. We also obtain a complete characterization of relative expressiveness and succinctness of modal logics defined in this framework.
2. "Several-Step" operators address worlds that can be reached in arbitrarily many steps. For our result, it suffices to study operators defined by a language L over $\{1, \ldots, n\}$: The formula $\Box_L \varphi$, evaluated in a world w, requires φ to be true in all worlds w' that can be reached from w on a path whose labels form a word in L. We show that using *alternation languages* suffices to get arbitrarily complex expressiveness- and succinctness relationships.

Most of our proofs use formula size games for modal logic as introduced in [FvdHIK13], based on Adler-Immerman games defined in [AI03]. These techniques allow to use games similar to Ehrenfeucht–Fraïssé-games to obtain lower bounds on formula size instead of quantifier depth. We adept these games to our generalized settings in the natural way. To the two techniques for establishing lower bounds in Adler-Immerman games mentioned in [FvdHIK13] (namely, *Diverging Pairs* and *Weight Function*), we add a third technique, which is based on a pigeon-hole principle argument.

The paper is structured as follows: Section 1 contains the classical definitions of syntax and semantics for modal logics. Section 2 contains our main results as outlined above. These results are based on a more detailed study of expressiveness and succinctness in the two settings we use, which forms the remainder of the paper: After reviewing formula size games for modal logic introduced in [FvdHIK13] in Section 3, we present our results on "One-Step" and "Several-Step" operators in Sections 4 and 5, respectively. We conclude in Section 6. All proofs can be found in the appendix.

1 Preliminaries

We fix an infinite set V of propositional variables. A Kripke model with n modalities is a tuple $M = (W, R_1, \ldots, R_n, \Pi)$, where W is a non-empty set of worlds

and for each $i \in \{1, \ldots, n\}$, R_i is a subset of $W \times W$, and $\Pi \colon P \to 2^W$ is a propositional assignment. We often simply write $w \in M$ for a world $w \in W$, and $S \subseteq M$ for $S \subseteq W$. A *pointed model* is a pair (M, w) where M is a Kripke model and w is a world of M. We also call w the *root* of the pointed model.

The semantics of a modal operator is characterized by the set of worlds that the operator "adresses" when evaluated in a pointed model. We formalize this as a *successor selection function*, which is a function O that for each pointed model (M, w) with n modalities returns a set a set $S \subseteq M$. We call n the *arity* of O. (Our definition does not rule out mal-formed successor selection functions that do not respect the relational character of modal logic, however all operators we study in this paper are "well-behaved" in that sense.)

A successor selection function O naturally defines a modal operator $\Box_O$, by replacing the "all successors" of the classical $\Box$-operator with "all worlds returned by O" (see the formal semantics below). Each set of successor selection functions defines a modal logic as follows (we identify a modal logic with the set of its formulas, as the satisfaction definition will always be standard).

Definition 1.1. *Let $\mathcal{O}$ be a set of successor selection functions. The modal logic $\mathsf{ML}^{\mathcal{O}}$ is generated by the following grammar:*

$$\varphi := p \mid \neg\varphi \mid \varphi \vee \varphi \mid \Box_O \varphi,$$

where $p \in P$ and $O \in \mathcal{O}$.

The *size* of a modal formula φ, denoted $|\varphi|$, is the number of nodes in its tree representation. The semantics definition of $\mathsf{ML}^{\mathcal{O}}$ is the natural one:

Definition 1.2. *Let φ be an $\mathsf{ML}^{\mathcal{O}}$-formula, and let (M, w) be a pointed model, where $M = (W, R_1, \ldots, R_n, \Pi)$. We define when φ is* satisfied *in w, written as $M, w \models \varphi$:*

- *$M, w \models p$ if and only if $w \in \Pi(p)$,*
- *$M, w \models \varphi \vee \psi$ if and only if $M, w \models \varphi$ or $M, w \models \psi$,*
- *$M, w \models \neg\varphi$ if and only if $M, w \not\models \varphi$*
- *$M, w \models \Box_O \varphi$ if $M, w' \models \varphi$ for all $w' \in O(M, w)$.*

For a set $\mathbb{M}$ of pointed models and a modal formula φ, we write $\mathbb{M} \models \varphi$ if $M, w \models \varphi$ for each $(M, w) \in \mathbb{M}$. Formulas φ and ψ are *equivalent* if for every pointed model (M, w), we have that $M, w \models \varphi$ if and only if $M, w \models \psi$.

We now define when one modal logic is more expressive or succinct than another. We only state these definitions as far as relevant for this paper, and refer the reader to [FvdHIK13] for an in-depth discussion of these notions.

Definition 1.3. *Let $\mathcal{O}_1$ and $\mathcal{O}_2$ be sets of successor selection functions.*

- $\mathsf{ML}^{\mathcal{O}_2}$ *is* at least as expressive *(at least as succinct) as $\mathsf{ML}^{\mathcal{O}_1}$, written as $\mathsf{ML}^{\mathcal{O}_1} \leq_{expr} \mathsf{ML}^{\mathcal{O}_2}$ ($\mathsf{ML}^{\mathcal{O}_1} \leq_{poly} \mathsf{ML}^{\mathcal{O}_2}$), if for every $\mathsf{ML}^{\mathcal{O}_1}$-formula φ, there is an equivalent $\mathsf{ML}^{\mathcal{O}_2}$-formula ψ (and $|\psi| \leq p(|\psi|)$ for a fixed polynomial p). If $\mathsf{ML}^{\mathcal{O}_1} \leq_{expr} \mathsf{ML}^{\mathcal{O}_2}$ and $\mathsf{ML}^{\mathcal{O}_2} \leq_{expr} \mathsf{ML}^{\mathcal{O}_1}$, then $\mathsf{ML}^{\mathcal{O}_1}$ and $\mathsf{ML}^{\mathcal{O}_2}$ are* equally expressive.

– *If* $\mathsf{ML}^{\mathcal{O}_1}$ *and* $\mathsf{ML}^{\mathcal{O}_2}$ *are equally expressive, then* $\mathsf{ML}^{\mathcal{O}_1}$ *is* exponentially more succinct *than* $\mathsf{ML}^{\mathcal{O}_2}$*, if there is a sequence* $(\varphi_i)_{i\in\mathbb{N}}$ *of* $\mathsf{ML}^{\mathcal{O}_1}$*-formulas such that* $|\varphi_i|$ *is linear in* i*, and there is some* $c > 1$ *such that for each* i*, each* $\mathsf{ML}^{\mathcal{O}_2}$*-formula equivalent to* φ_i *has size at least* c^i*.*

Clearly, $\mathsf{ML}^{\mathcal{O}_1} \not\leq_{expr} \mathsf{ML}^{\mathcal{O}_2}$ does not imply that $\mathsf{ML}^{\mathcal{O}_1}$ is more expressive than $\mathsf{ML}^{\mathcal{O}_2}$, since $\mathsf{ML}^{\mathcal{O}_2} \not\leq_{expr} \mathsf{ML}^{\mathcal{O}_1}$ can hold simultaneously. If $\mathsf{ML}^{\mathcal{O}_1}$ is exponentially more succinct than $\mathsf{ML}^{\mathcal{O}_2}$, then an exponential blow-up in the translation from $\mathsf{ML}^{\mathcal{O}_1}$ to $\mathsf{ML}^{\mathcal{O}_2}$ cannot always be avoided, on the other hand, if $\mathsf{ML}^{\mathcal{O}_1} \leq_{poly} \mathsf{ML}^{\mathcal{O}_2}$, then every $\mathsf{ML}^{\mathcal{O}_1}$-formula can be *succinctly* rewritten into a $\mathsf{ML}^{\mathcal{O}_2}$-formula.

One needs to be careful when proving succinctness result via a complexity argument: Unless PSPACE = P, there is no polynomial-time algorithm converting every closed QBF-formula into a constant formula. However, since each closed QBF-formula is equivalent to either *true* or *false*, the class of closed QBF-formulas is certainly not more succinct than the class of constant formulas. We do not discuss these issues further, since in this paper we will always have that if we compare $\mathcal{L}_1$ and $\mathcal{L}_2$ that are equally expressive, then either $\mathcal{L}_1 \leq_{poly} \mathcal{L}_2$ and the translation can be computed by a polynomial-time algorithm, or $\mathcal{L}_1$ is exponentially more succinct than $\mathcal{L}_2$ in the above, strict sense.

2 Main Results

We prove that the expressiveness- and succinctness relationships between modal logics can be as complex as any partial order. We show versions of this result in two settings: 1. For logics defined by successor selection functions O such that whether $w' \in O(M, w)$ only depends on whether $(w, w') \in R_i$ for each accessibility relation R_i, 2. for logics defined by successor selection functions considering paths of arbitrary (finite) length in the model.

A simple counting argument shows that in the first setting, there is only a finite number of different modal operators, hence for these operators we show that the relationships can be as complex as any finite partial order. In the second setting, we then obtain relationships as complex as any countable partial order.

2.1 Single Step Operators

In order to prove that the relationships between different modal logics can be arbitrarily complex, we first define a large class of modal logics. All of our logics will be extensions of the classical multi-modal logic ML_n. As a starting point, consider the following modal operators (see also [vdHI14]): For a set $I \subseteq \{1, \ldots, n\}$,

– $[\forall_I]\varphi$ is true in w if φ is true in all w' such that $(w, w') \in R_i$ for some $i \in I$.
– $[\cap_I]\varphi$ is true in w if φ is true in all w' such that $(w, w') \in R_i$ for all $i \in I$.

The first of these operators can be expressed with standard multimodal logic, since $[\forall_I]\varphi$ is equivalent to $\wedge_{i\in I}\Box_i\varphi$. The second one cannot be expressed, since in

the standard modal language, we cannot "address" a world w' based on whether, for example, $(w, w') \in R_1$ *and* $(w, w') \in R_2$ both hold at the same time. In this section, we consider successor selection functions O that can do exactly this: Whether $w' \in O(M, w)$ depends on the R_i-relationships between w and w_i for all i *simultaneously*.

More precisely, we consider n-ary successor selection functions O for which the question whether $w' \in O(M, w)$ is described as a Boolean combination of whether $(w, w') \in R_i$ for each relevant i. Such an O is is characterized by a Boolean function $f\colon \{0,1\}^n \to \{0,1\}$ as follows: For worlds w, w' of a model M, we say that w' is an f-successor of w if $f(r_1, \ldots, r_n) = 1$, where $r_i = 1$ if $(w, w') \in R_i$, and $r_i = 0$ otherwise. Then f yields a successor selection function in the obvious way:

Definition 2.1. *Let $f\colon \{0,1\}^n \to \{0,1\}$. Then for a pointed model (M, w) with $M = (W, R_1, \ldots, R_n, \Pi)$, $O_f(M, w) = \{w' \mid w'$ is an f-successor of $w\}$.*

We often identify a Booelan function f and the successor selection function O_f defined by f. Hence for a set $\mathcal{F}$ of Boolean functions, we use $\mathsf{ML}^{\mathcal{F}}$ to denote the modal logic $\mathsf{ML}^{\{O_f \mid f \in \mathcal{F}\}}$, write $\Box_f$ instead of $\Box_{O_f}$, etc. The usual multi-model logic with n modalities is obtained as $\mathsf{ML}_n = \mathsf{ML}^{\{r_1, \ldots, r_n\}}$ (we identify a Boolean function with the propositional formula over the variables $\{r_1, \ldots, r_n\}$ representing it, hence using the notation above, $\Box_{r_i}\varphi$ is equivalent to $\Box_i\varphi$). As an example, the above operator $[\forall_I]$ corresponds to the successor selection function $O_{\forall_I}(r_1, \ldots, r_n) = \vee_{i \in I} r_i$: It addresses all worlds w' that are an i-successor of w for *some* $i \in I$. The operator $[\cap_I]$ similarly corresponds to the successor selection function $O_{\cap_I}(r_1, \ldots, r_n) = \wedge_{i \in I} r_i$, as it selects all worlds w' such that $(w, w') \in R_i$ for *all* $i \in I$.

We now state our main result for modal logics of the form $\mathsf{ML}^{\mathcal{F}}$: The expressiveness and succinctness relationships between logics $\mathsf{ML}^{\mathcal{F}}$ can be as complex as any finite partial order.

Theorem 2.2. *Let S be a finite set, and let $\leq_S$ be a partial order on S. Then there exist families of sets of $\lceil \log_2(|S| + 1) \rceil$-ary Boolean functions $(\mathcal{F}_s)_{s \in S}$ and $(\mathcal{G}_s)_{s \in S}$ such that for each $s, t \in S$, the following holds:*

1. *$\mathsf{ML}^{\mathcal{F}_s} \leq_{expr} \mathsf{ML}^{\mathcal{F}_t}$ if and only if $\mathsf{ML}^{\mathcal{G}_s} \leq_{poly} \mathsf{ML}^{\mathcal{G}_t}$ if and only if $s \leq_S t$.*
2. *All logics $\mathsf{ML}^{\mathcal{G}_s}$ are equally expressive, and if $s \not\leq_S t$, then $\mathsf{ML}^{\mathcal{G}_s}$ is exponentially more succinct that $\mathsf{ML}^{\mathcal{G}_t}$.*

In particular, if s and t are not comparable with respect to $\leq_S$, then $\mathsf{ML}^{\mathcal{F}_s}$ is exponentially more succinct than $\mathsf{ML}^{\mathcal{F}_t}$ and vice versa, and there are formulas expressible in $\mathsf{ML}^{\mathcal{G}_s}$ but not in $\mathsf{ML}^{\mathcal{G}_t}$ and vice versa. To prove Theorem 2.2, we study the expressivity- and succinctness relationship between $\mathsf{ML}^{\mathcal{F}}$ and $\mathsf{ML}^{\mathcal{G}}$ for different sets $\mathcal{F}$ and $\mathcal{G}$ in detail, and obtain a complete characterization that for each $\mathcal{F}$ and $\mathcal{G}$ determines the precise relationship between $\mathsf{ML}^{\mathcal{F}}$ and $\mathsf{ML}^{\mathcal{G}}$ in terms of $\leq_{expr}$, $\leq_{poly}$, and exponential succinctness. These results can be found in Section 4.

2.2 Arbitrary Step Operators

In this section, we obtain an "infinite version" of Theorem 2.2. As argued above, for a fixed arity n, there is only a finite number of modal logics of the form $\mathsf{ML}^{\mathcal{F}}$ on Kripke models with n modalities. Hence we consider logics outside of the above framework, i.e., successor selection functions O where whether $w' \in O(M, w)$ does not only depend on whether $(w, w') \in R_i$ for each i, but also on longer paths in the model. Natural functions of this form are, e.g., ones returning all worlds reachable on a path of a certain maximal length, or on a path of arbitrary length (which allows to express the transitive closure of the accessibility relations). For our result, it suffices to consider operators of a simple structure, which for a Kripke model with n modalities are given by languages over $\{1, \dots, n\}$. For a word $s = s_1 \dots s_l \in \{1, \dots, n\}^*$, we say that a world w' is an s-successor of a world w in a model M if there are worlds $w = w_0, w_1, \dots, w_l = w'$ such that for each $i \in \{1, \dots, l\}$, we have that $(w_{i-1}, w_i) \in R_{s_i}$. In this case we say that there is an s-path from w to w' in M, and refer to the s_i as the *labels* of this path. (We omit the model when clear from the context).

A language $L \subseteq \{1, \dots, n\}^*$ defines the successor selection function $O_L(M, w) = \{w' \in M \mid w' \text{ is an } s\text{-successor of } w \text{ for some } s \in L\}$. Again, we identify a language L and the successor selection function O_L, e.g., we write $\Box_L$ instead of $\Box_{O_L}$, and $\mathsf{ML}^{\mathcal{L}}$ for $\mathsf{ML}^{\{O_L \mid L \in \mathcal{L}\}}$, etc. The usual multi-model logic with n modalities is obtained as $\mathsf{ML}_n = \mathsf{ML}^{\{\{1\}, \dots, \{n\}\}}$.

In the sequel, we only consider finite languages. Clearly, for a set $\mathcal{L}$ of finite languages, every $\mathsf{ML}^{\mathcal{L}}$-formula is equivalent to some ML_n-formula, since $\Box_L \varphi$ is equivalent to $\bigwedge_{s=s_1 s_2 \dots s_k \in L} \Box_{s_1} \Box_{s_2} \dots \Box_{s_k} \varphi$ for a finite language L.

Our main result for logics of the form $\mathsf{ML}^{\mathcal{L}}$ is an "infinite version" of Theorem 2.2: The succinctness- and expressiveness- relationships between modal logics of the form $\mathsf{ML}^{\mathcal{L}}$ can be as complex as any *countable* partial order. For the result, it suffices to consider the bimodal case, i.e., models (W, R_1, R_2, Π) with two accessibility relations, and languages over the alphabet $\{1, 2\}$.

Theorem 2.3. *Let S be a countable set, and let $\leq_S$ be a partial order on S. Then there exist families of languages $(\mathcal{L}_s)_{s \in S}$ and $(\mathcal{K}_s)_{s \in S}$ over the alphabet $\{1, 2\}$ such that for each $s, t \in S$, the following holds:*

1. *$\mathsf{ML}^{\mathcal{K}_s} \leq_{expr} \mathsf{ML}^{\mathcal{K}_t}$ if and only if $\mathsf{ML}^{\mathcal{L}_s} \leq_{poly} \mathsf{ML}^{\mathcal{L}_t}$ if and only if $s \leq_S t$.*
2. *All logics $\mathsf{ML}^{\mathcal{L}_s}$ are equally expressive, and if $s \not\leq_S t$, then $\mathsf{ML}^{\mathcal{L}_s}$ is exponentially more succinct that $\mathsf{ML}^{\mathcal{L}_t}$.*

We will give an overview of the proof in Section 5.

3 Formula Size Games

Our succinctness proofs use modal formula size games introduced in [FvdHIK13] building on Adler-Immerman games [AI03]. We review these games in Section 3.1, and state a variation of their formula-size theorem in Section 3.2. In Section 3.3, we introduce the pigeonhole-technique to prove lower bounds on the size of game trees (which then translate to lower bounds on formula size).

3.1 Game Trees

The following definition is taken from [FvdHIK13], except for the straightforward extension to $\Box_O$-moves. A game tree represents a formula, where each node v corresponds to a subformula v_φ in the natural way. A node v has labels of two kinds: The first label is of the form $\langle \mathbb{A} \circ \mathbb{B} \rangle$, where $\mathbb{A}$ and $\mathbb{B}$ are classes of pointed models such that $\mathbb{A} \models v_\varphi$ and $\mathbb{B} \models \neg v_\varphi$. The second label contains the outmost operator of the formula v_φ. We simply refer to both labels as "label," it will always be clear whether we refer to the models or the operators. In the following definition, the goal of the single player "Spoiler" is to find a formula that is true on all models in $\mathbb{A}$, and false on all models in $\mathbb{B}$. Successful plays of Spoiler (called *closed game trees*) directly correspond to such formulas.

Definition 3.1 ([FvdHIK13]). *The formula-size game for a set $\mathcal{O}$ of successor selection functions (FSG($\mathcal{O}$)) on two sets of pointed models $\mathbb{A}$ and $\mathbb{B}$ is played as follows: The game begins with a tree containing only the root labelled $\langle \mathbb{A} \circ \mathbb{B} \rangle$. In each move of the game, the player (Spoiler) chooses a leaf that is labelled $\langle \mathbb{C} \circ \mathbb{D} \rangle$ for classes $\mathbb{C}$ and $\mathbb{D}$ of pointed models and not labelled with a variable, and plays one of the following moves:*

atomic move *Spoiler labels the leaf p for a propositional variable p such that $\mathbb{C} \models p$ and $\mathbb{D} \models \neg p$.*

not move *Spoiler labels the leaf with $\neg$ and adds a new leaf $\langle \mathbb{D} \circ \mathbb{C} \rangle$ as successor.*

or move *Spoiler labels the leaf with $\vee$ and chooses two subsets $\mathbb{C}_1, \mathbb{C}_2 \subseteq \mathbb{C}$ with $\mathbb{C} = \mathbb{C}_1 \cup \mathbb{C}_2$, then adds successor nodes labelled $\langle \mathbb{C}_1 \circ \mathbb{D} \rangle$ and $\langle \mathbb{C}_2 \circ \mathbb{D} \rangle$.*

$\Box_O$-move *Spoiler labels the leaf with $\Box_O$ for some $O \in \mathcal{O}$ and chooses a set $\mathbb{D}_1$ such that for each $(M, w) \in \mathbb{D}$, there is some $(M, w') \in \mathbb{D}_1$ with $w' \in O(M, w)$. A new successor node $\langle \mathbb{C}_1 \circ \mathbb{D}_1 \rangle$ is added to the tree, where $\mathbb{C}_1 = \{(M, w') \mid (M, w) \in \mathbb{C}, w' \in O(M, w)\}$.*

A game tree is closed *if all of its leafs are labelled with variables.*

By definition, Spoiler cannot play an $\Box_O$-move on a node $\langle \mathbb{C} \circ \mathbb{D} \rangle$ if there is some $(M, w) \in \mathbb{D}$ with $O(M, w) = \emptyset$ (this reflects that $M, w \models \Box_O \varphi$ for all φ in this case). The set $\mathcal{T}_\mathcal{O}(\langle \mathbb{A} \circ \mathbb{B} \rangle)$ contains all closed game trees of FSG($\mathcal{O}$) with a root labelled $\langle \mathbb{A} \circ \mathbb{B} \rangle$. Spoiler wins the FSG($\mathcal{O}$) starting at $\langle \mathbb{A} \circ \mathbb{B} \rangle$ in n moves if there is some $T \in \mathcal{T}_\mathcal{O}(\langle \mathbb{A} \circ \mathbb{B} \rangle)$ with exactly n nodes. We usually only write $\mathcal{T}(\langle \mathbb{A} \circ \mathbb{B} \rangle)$ instead of $\mathcal{T}_\mathcal{O}(\langle \mathbb{A} \circ \mathbb{B} \rangle)$ if the set $\mathcal{O}$ is clear from the context.

3.2 Formula Size Game Theorem

The proof of Thoerem 1 from [FvdHIK13] can be generalized in a straightforward way to give the following result (for completeness, we give the complete proof in Appendix B.1.1).

Theorem 3.2. [FvdHIK13] *Spoiler wins the FSG($\mathcal{O}$) starting with $\langle \mathbb{A} \circ \mathbb{B} \rangle$ in k moves if and only if there is a formula $\varphi \in \mathsf{ML}^\mathcal{O}$ with $|\varphi| = k$ such that $\mathbb{A} \models \varphi$ and $\mathbb{B} \models \neg\varphi$.*

3.3 Pigeonhole Principle Technique

The Formula Size Theorem (Theorem 3.2) allows to prove lower bounds on a $\mathsf{ML}^{\mathcal{O}}$-formula φ by showing a lower bound on the smallest game tree in $\mathcal{T}_{\mathcal{O}}(\langle \mathbb{A} \circ \mathbb{B} \rangle)$, where $\mathbb{A} \models \varphi$ and $\models \neg\varphi$. However, proving a lower bound for game trees is a nontrivial task itself. In [FvdHIK13], two techniques for proving such a lower bound are mentioned, namely, *Diverging Pairs* and using a *Weight Function*. For our result, we use a different technique, which is based on a Pigeonhole-like counting argument. The idea is to show that each branch of a formula can only "cover" a certain number c of models from $\mathbb{A}$. From this it then easily follows that the formula must have at least $\frac{|\mathbb{A}|}{c}$ nodes.

The result uses that formula size games allow the classes of models "covered" by each branch of a closed tree (corresponding to a formula) to be simply read off the labels of the leaf of the branch. For a tree $T \in \mathcal{T}(\langle \mathbb{A} \circ \mathbb{B} \rangle)$ and a node v of T labelled $\langle \mathbb{C} \circ \mathbb{D} \rangle$, we say that $\mathbb{C}$ ($\mathbb{D}$) is the class *corresponding to* $\mathbb{A}$, if there is an even (odd) number of negations on the path from T's root to v, and the class *corresponding to* $\mathbb{B}$ otherwise.

Theorem 3.3. *Let $\mathcal{O}$ be a set of successor selection functions. Let φ be a formula, let $\mathbb{A}$ and $\mathbb{B}$ be sets of pointed models such that $\mathbb{A} \models \varphi$ and $\mathbb{B} \models \neg\varphi$. If for every nontrivial leaf u of every closed game tree $T_{\mathcal{O}} \in \mathcal{T}(\langle \mathbb{A} \circ \mathbb{B} \rangle)$, the class of models corresponding to $\mathbb{A}$ ($\mathbb{B}$) has size at most c, then every $\mathsf{ML}^{\mathcal{O}}$-formula equivalent to φ has size at least $\frac{|\mathbb{A}|}{c}$ ($\frac{|\mathbb{B}|}{c}$).*

4 Succinctness and Expressiveness for Single-Step Operators

In this section, we study the expressiveness- and succinctness relationships between logics of the form $\mathsf{ML}^{\mathcal{F}}$ for classes $\mathcal{F}$ of Boolean functions. In particular, these results allow us to prove the above Theorem 2.2. We first consider expressiveness. The following result completely answers the question in which case $\mathsf{ML}^{\mathcal{G}} \leq_{expr} \mathsf{ML}^{\mathcal{F}}$ holds:

Theorem 4.1. *Let $\mathcal{F}$ and $\mathcal{G}$ be sets of n-ary Boolean functions. Then the following are equivalent:*

1. $\mathsf{ML}^{\mathcal{G}} \leq_{expr} \mathsf{ML}^{\mathcal{F}}$,
2. *for each $g \in \mathcal{G}$, there is a set $S \subseteq \mathcal{F}$ such that $g \equiv \bigvee_{f \in S} f$.*

For example, the theorem implies the result mentioned in Section 2.1 that $[\forall_I]$ can be expressed with the standard operators $\Box_1$ and $\Box_2$, but $[\cap_I]$ cannot (recall that $[\forall_I]$ corresponds to $\vee_{i \in I} r_i$, and $[\cap_I]$ to $\wedge_{i \in I} r_i$).

Theorem 4.1 is proved using standard bisimulation techniques (see Appendix B.2.1), which show that a specific formula cannot be expressed in a logic $\mathsf{ML}^{\mathcal{F}}$. We now consider succinctness. The following theorem says that, given sets $\mathcal{F}$ and $\mathcal{G}$ of Boolean functions such that $\mathsf{ML}^{\mathcal{F}}$ and $\mathsf{ML}^{\mathcal{G}}$ are equally expressive, $\mathsf{ML}^{\mathcal{G}}$ is *always* exponentially more succinct than $\mathsf{ML}^{\mathcal{F}}$, except for the trivial case when $\mathcal{G} \subseteq \mathcal{F}$.

The proof of the theorem indeed shows the slightly stronger result that even if $\mathsf{ML}^{\mathcal{F}}$ and $\mathsf{ML}^{\mathcal{G}}$ are not equally expressive, but $\mathcal{G}$ contains a function that is a disjunction of functions in $\mathcal{F}$ but is not an element of $\mathcal{F}$ itself (and hence, due to Theorem 4.1, $\Box_g$ is not expressible in $\mathsf{ML}^{\mathcal{F}}$), then $\mathsf{ML}^{\mathcal{G}}$ is exponentially more succinct than $\mathsf{ML}^{\mathcal{F}}$ (with a slightly more general definition of this notion that also covers modal languages with different expressive power). This implies that the relation $\leq_{poly}$ restricted to logics of the form $\mathsf{ML}^{\mathcal{F}}$ is antisymmetric, and hence a partial order.

A special case of our result was shown in [FvdHIK13], where the authors prove that the logic $[\forall_{1,2}]\mathsf{ML}$ is exponentially more succinct than ML_2. Using our notation, these logics are $\mathsf{ML}^{\mathcal{G}}$ and $\mathsf{ML}^{\mathcal{F}}$ with $\mathcal{F} = \{r_1, r_2\}$ and $\mathcal{G} = \{r_1, r_2, r_1 \vee r_2\}$.

Theorem 4.2. *Let $\mathcal{F}$ and $\mathcal{G}$ be sets of Boolean functions such that $\mathsf{ML}^{\mathcal{F}}$ and $\mathsf{ML}^{\mathcal{G}}$ are equally expressive and $\mathcal{G} \not\subseteq \mathcal{F}$. Then $\mathsf{ML}^{\mathcal{G}}$ is exponentially more succinct than $\mathcal{F}$.*

The proof of Theorem 4.2 uses an extension of the technique used to prove the above-mentioned result in [FvdHIK13]. The main additions we make to their construction are "false paths" in the models that stop Spoiler from using operators of the form $\Box_f$ where f is not one of the functions from $\mathcal{F}$ appearing in the disjunctive definition of g, and a generalization of edges labelled with 1 and 2 to edges labelled with appropriate Boolean combinations of the involved modalities. Finally, instead of the diverging pairs technique, we use the pigeonhole technique to prove the lower bound on the game tree size. The proof can be found in Appendix B.2.2 (In Section 5, we give a more detailed presentation of an application of the pigeonhole technique.)

5 Succinctness and Expressiveness for Arbitrary-Step Operators

In this section, we give an overview of the proof of Theorem 2.3. In particular, we define the sets of languages $\mathcal{L}_s$ mentioned in the statement of the theorem as sets of *alternation languages* (Section 5.1). We then study the relationships between $\mathsf{ML}^{\mathcal{L}_1}$ and $\mathsf{ML}^{\mathcal{L}_2}$ for sets $\mathcal{L}_1$ and $\mathcal{L}_2$ of alternation languages in detail. Due to the page limit, we only give the construction (Section 5.2) and state its main technical properties (Section 5.3), and the consequences for expressiveness (Section 5.4) and succinctness (Section 5.5). The technical proofs are deferred to Appendix B.3.

To prove Theorem 2.3, it is enough to consider bimodal logics, i.e., models with two accessibility relations and thus languages over the alphabet $\{1, 2\}$. We therefore only consider this case in the remainder of this section.

5.1 Alternation Languages

Let $\ell \geq 1$ be a natural number. A word $s = s_1 \dots s_\ell \in \{1, 2\}^\ell$ is *alternating* if for each $i \in \{1, \dots, \ell - 1\}$, $s_i \neq s_{i+1}$. There are exactly two alternating words

of length ℓ, namely a_1^ℓ, starting with 1, and a_2^ℓ, starting with 2. The *alternation language of length* ℓ, denoted with A_ℓ, is the set $\{a_1^\ell, a_2^\ell\}$.

Following the definitions in Section 2.2, the language A_ℓ defines the modal operator $\Box_{A_\ell}$, where $\Box_{A_\ell}\varphi$ requires φ to be true in all worlds reachable on a path whose labels form an alternating word of length ℓ. This operator is natural in an epistemic setting, where it can be read as "A knows that B knows that A knows that B knows ..." and vice versa, to the ℓ-th degree.

The iterated application of the operator $\Box_{A_\ell}$, denoted as usual with $\Box^i_{A_\ell}$, addresses all worlds accessible on a path whose labels form a sequence of i words from A_ℓ. To be able to address the specific alternating words in this sequence, we extend the notation a_1^ℓ and a_2^ℓ above: For a word $s = s_1 \dots s_i \in \{1,2\}^*$, with a_s^ℓ we denote the word $a_{s_1}^\ell a_{s_2}^\ell \dots a_{s_i}^\ell$, i.e., the word consisting of i alternating words of length ℓ, where the j-th of these words starts with s_j.

For a set $I \subseteq \mathbb{N}$, let $\mathsf{ML}^{A(I)}$ denote the logic $\mathsf{ML}^{\{A_\ell \mid \ell \in I\}}$, and let $\mathsf{ML}^{A(I)^+}$ denote the logic $\mathsf{ML}^{\{A_\ell \mid \ell \in I\} \cup \{\{1\},\{2\}\}}$. Hence in the logic $\mathsf{ML}^{A(I)}$, all operators $\Box_{A_\ell}$ with $\ell \in I$ are allowed, the logic $\mathsf{ML}^{A(I)^+}$ additionally allows the classical operators $\Box_1$ and $\Box_2$. Since all logics $\mathsf{ML}^{A(I)^+}$ contain $\Box_1$ and $\Box_2$, and all involved languages are finite, all $\mathsf{ML}^{A(I)^+}$ have the same expressive power, namely that of classical bimodal logic. This is not true for the logics $\mathsf{ML}^{A(I)}$, as we will see below.

Our main result on alternation languages states that if I_1 is not a subset of I_2, then $\mathsf{ML}^{A(I_1)^+}$ is exponentially more succinct than $\mathsf{ML}^{A(I_2)^+}$, and $\mathsf{ML}^{A(I_1)}$ contains formulas that are not expressible in $\mathsf{ML}^{A(I_2)}$. This result is the key step to proving Theorem 2.3.

Theorem 5.1. *Let $I_1, I_2 \subseteq \mathbb{N}$ with $I_1 \nsubseteq I_2$. Then*

1. *$\mathsf{ML}^{A(I_1)^+}$ is exponentially more succinct than $\mathsf{ML}^{A(I_2)^+}$, and*
2. *there is an $\mathsf{ML}^{A(I_1)}$-formula for which there is no equivalent $\mathsf{ML}^{A(I_2)}$-formula.*

Note that, in contrast to the situation for logics of the form $\mathsf{ML}^{\mathcal{F}}$ for a set of Boolean functions $\mathcal{F}$ (see Section 4), we do not get the corresponding result that the relation $\leq_{poly}$, restricted to logics of the form $\mathsf{ML}^{\mathcal{L}}$ is antisymmetric. The reason for this is that by adding an operator $\Box_L$ for a *singleton* language to a logic containing both classical operators $\Box_1$ and $\Box_2$ changes neither expressiveness nor succinctness of the logic; hence an arbitrary number of logics equivalent to ML_n in expressiveness and succinctness can be defined in this way.

5.2 Model Construction

Our main result about alternation languages, and the main ingredient to the proof of Theorem 5.1, is that if $\ell \notin I$, then every $\mathsf{ML}^{A(I)^+}$-formula ψ equivalent to $\Box^i_{A_\ell} p$ is exponentially large (in i), and there is no $\mathsf{ML}^{A(I)}$-formula equivalent to $\Box_{A_\ell} p$. In the following discussion, we focus on the more involved succinctness result.

We start by defining the models on which we play the formula size game: For each i and $\ell \geq 1$, we define classes of pointed models $\mathbb{A}_i^{\ell,*}$ and $\mathbb{B}_i^{\ell,*}$ such that $\mathbb{A}_i^{\ell,*} \models \Box^i_{A_\ell} p$, and $\mathbb{B}_i^{\ell,*} \models \neg \Box^i_{A_\ell} p$.

These classes are defined in three steps:

1. For each $i \in \mathbb{N}$, we define a *base model* $\mathbb{A}_i^\ell$, and for each $s \in \{1,2\}^i$, a *base model* $\mathbb{B}_s^\ell$.
2. We then extend the models $\mathbb{A}_i^\ell$ and $\mathbb{B}_s^\ell$ to $\mathbb{A}_i^{\ell,*}$ and $\mathbb{B}_s^{\ell,*}$. The extension consists of adding a "trap" to the models which the Spoiler candidate may never choose as successors in the $\Box_O$-step of the formula size games. This
3. for each $i \in \mathbb{N}$, we define the above-mentioned classes. Since for each such i, there is only a single model $\mathbb{A}_i^{\ell,*}$, we wimply identify the model $\mathbb{A}_i^{\ell,*}$ with the class $\left\{\mathbb{A}_i^{\ell,*}\right\}$ and write only $\mathbb{A}_i^{\ell,*}$ for this class. On the other hand, the class $\mathbb{B}_i^{\ell,*}$ contains all models $\mathbb{B}_s^{\ell,*}$ with $s \in \{1,2\}^i$.

The main idea of the definition of our models is to ensure that each branch of each tree $T \in \mathcal{T}_\mathcal{O}(\langle \mathbb{A}_i^{\ell,*} \circ \mathbb{B}_i^{\ell,*} \rangle)$ corresponding to a smallest formula ψ as above can only cover a restricted number of models (namely, at most $2^{\frac{i}{2}}$ models). This allows us to use the pigeonhole technique (Theorem 3.3) to prove our succinctness result.

For $i, \ell \in \mathbb{N}$, let $s = s_1 \dots s_i \in \{1,2\}^i$. The main idea of the following models is that in order to "cover" all models $\mathbb{B}_s^{\ell,*}$, Spoiler needs to exhibit, for each of them, an A_ℓ^i-successor of the root in which the variable p is false. The "traps" in the models $\mathbb{B}_s^{\ell,*}$ ensure that the only path for which this is true is the word a_s^ℓ. Therefore, Spoiler needs to play a strategy that covers at least all strings of the form a_s^ℓ. The model $\mathbb{A}_i^{\ell,*}$ forces Spoiler to only cover strings of this form. Hence Spoiler must cover *exactly* all strings a_s^ℓ where $s \in \{1,2\}^i$, which cannot be done succinctly without the operator $\Box_{A_\ell}$. The "base models" models $\mathbb{A}_i^\ell$ and $\mathbb{B}_s^\ell$ are defined as in the following picture:

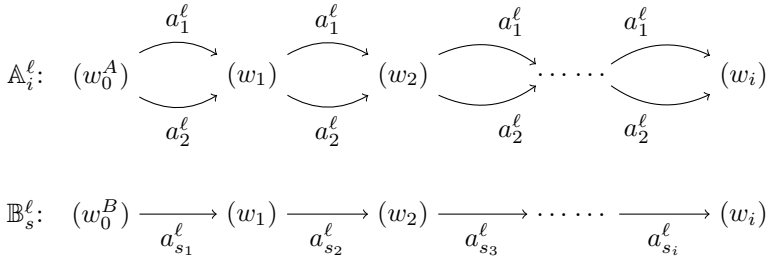

An edge labelled a_j^ℓ for $j \in \{1,2\}$ between w_m and w_{m+1} indicates that w_{m+1} is an a_j^ℓ-successor of w_m. This is achieved by intermediate worlds (not shown in the picture) $w_m = u_0, \dots, u_\ell = w_{m+1}$ such that $(u_{p-1}, u_p) \in R_{a_j^\ell[p]}$ for each relevant p. The propositional variable p is false in all worlds, except for the world w_i of $\mathbb{A}_i^\ell$. The root of $\mathbb{A}_i^\ell$ ($\mathbb{B}_s^\ell$) is w_0^A (w_0^B). By construction, each world u in $\mathbb{A}_i^\ell$ and $\mathbb{B}_s^\ell$ has a unique distance from the model's root. We denote this distance with $depth(u)$.

We obtain the "extended" models $\mathbb{A}_i^{\ell,*}$ and $\mathbb{B}_s^{\ell,*}$ from $\mathbb{A}_i^{\ell}$ and $\mathbb{B}_s^{\ell}$ as follows: To both models, we add a new node w_{trap}, which is a reflexive singleton (i.e., a world with an 1- and a 2-edge to itself). In $\mathbb{A}_i^{\ell,*}$, the variable p is false in w_{trap}, in $\mathbb{B}_s^{\ell,*}$, the variable is true in w_{trap}. For each world w of $\mathbb{A}_i^{\ell}$ ($\mathbb{B}_s^{\ell}$) that does not have a j-successor for some $j \in \{1,2\}$, we add a j-edge leading to the world w_{trap} of the respective model. These edges are "false paths," since the reflexive singleton of $\mathbb{A}_i^{\ell,*}$ does not allow Spoiler to prove that all relevant paths end in a world satisfying p, and the singleton in $\mathbb{B}_s^{\ell,*}$ does not allow Spoiler to find a path to a world where p is false. Hence these "false paths" are never taken in a closed game tree that corresponds to a minimal formula.

Our classes of models now contain all models constructed in the above way: For $i \geq 1$, we identify $\mathbb{A}_i^{\ell,*}$ with the singleton $\left\{\mathbb{A}_i^{\ell,*}\right\}$, and define $\mathbb{B}_i^{\ell,*} = \left\{\mathbb{B}_s^{\ell,*} \mid s \in \{1,2\}^i\right\}$. Then $\mathbb{A}_i^{\ell,*} \models \Box^i_{A_\ell} p$, and $\mathbb{B}_i^{\ell,*} \models \neg\Box^i_{A_\ell} p$:

- Every path made up of i alternating words of length ℓ starting at w_0^A in $\mathbb{A}_i^{\ell,*}$ leads to the world w_i of $\mathbb{A}_i^{\ell}$, where p is true (no such path ends in the reflexive singleton).
- For $s = s_1 \dots s_i \in \{1,2\}^i$, the world w_i, in $\mathbb{B}_s^{\ell,*}$, is an a_s^ℓ-successor of w_0^B and does not satisfy p.

5.3 Formula Size Games on our Models

We now state a few technical results on formula size games on closed game trees in $\mathcal{T}(\langle \mathbb{A}_i^{\ell,*} \circ \mathbb{B}_i^{\ell,*} \rangle)$. Essentially, these results say that Spoiler indeed needs to play a strategy as intended by the definition of our models. Recall that our goal is to show that if $\ell \notin I$, then every $\mathsf{ML}^{A(I)^+}$-formula ψ equivalent to $\Box^i_{A_\ell} p$ must be of exponential size. In the following, we fix a smallest such formula ψ, and consider the game tree that corresponds to the evaluation of ψ on the classes of models $\mathbb{A}_i^{\ell,*}$ and $\mathbb{B}_i^{\ell,*}$ in the following way: For a formula ψ and classes $\mathbb{A}$ and $\mathbb{B}$ of pointed models with $\mathbb{A} \models \psi$ and $\mathbb{B} \models \neg\psi$, let $T^\psi(\langle \mathbb{A} \circ \mathbb{B} \rangle)$ be the closed game tree obtained from following the strategy corresponding to ψ on the starting node $\langle \mathbb{A} \circ \mathbb{B} \rangle$. Clearly, $T^\psi(\langle \mathbb{A} \circ \mathbb{B} \rangle) \in \mathcal{T}_{\mathcal{O}}(\langle \mathbb{A} \circ \mathbb{B} \rangle)$ if $\mathcal{O}$ contains at least all succesor selection functions appearing in ψ.

We first show that the formula ψ indeed must indeed avoid the "traps" added to the models, as intended:

Lemma 5.2. *Let $I \subseteq \mathbb{N}$, let ψ be a minimal $\mathsf{ML}^{A(I)^+}$-formula equivalent to $\Box^i_{A_\ell} p$, let v be a node of $T^\psi(\langle \mathbb{A}_i^{\ell,*} \circ \mathbb{B}_i^{\ell,*} \rangle)$ labelled $\langle \mathbb{A} \circ \mathbb{B} \rangle$, and let $(M, w) \in \mathbb{A} \cup \mathbb{B}$. Then $w \neq w_{trap}$.*

Our next result is that in the formula ψ_i, operators $\Box_{A_{\ell'}}$ can only appear in depths that are multiples of ℓ. The proof uses that each path that is not a prefix of a word in $(A_\ell)^i$ leads to w_{trap} in the models from $\mathbb{A}_i^{\ell,*}$ and $\mathbb{B}_i^{\ell,*}$. For a node v of a tree T, with $labels_{\Box}(v)$, we denote the sequence of successor selection functions appearing in $\Box_\cdot$ operators on the path from T's root to v, excluding

the label of v itself. (We do not make T explicit in the notation, this will always be clear from the context, and again identify L and O_L for a language L).

We say that a language $\emptyset \neq L \subseteq \{1,2\}^*$ is *length*-uniform if there is some i such that $L \subseteq \{1,2\}^i$, i.e., all words in L have the same length. We denote this length i with $||L||$. Clearly, the class of length-uniform languages is closed under concatenation, and all languages L we consider in this section (the alternating languages A_ℓ and the languages $\{1\}$ and $\{2\}$) are length-uniform. For a node $v \in T^\psi(\langle \mathbb{A}_i^{\ell,*} \circ \mathbb{B}_i^{\ell,*} \rangle)$, and a string $s \in \{1,2\}^i$, we say that v *covers* s, if one of the classes of models with which v is labelled contains a model $(\mathbb{B}_s^{\ell,*}, w)$ for some $w \in \mathbb{B}_s^{\ell,*}$. As discussed before, we will show that each leaf $v \in T^\psi(\langle \mathbb{A}_i^{\ell,*} \circ \mathbb{B}_i^{\ell,*} \rangle)$ can only cover a restricted number of strings s.

Lemma 5.3. *Let ψ be a minimal $\mathsf{ML}^{A(I)^+}$-formula equivalent to $\Box^i_{A_\ell} p$. Let $v \in T^\psi(\langle \mathbb{A}_i^{\ell,*} \circ \mathbb{B}_i^{\ell,*} \rangle)$. Let $labels_\Box(v) = L_1 \dots L_{m-1} A_{\ell'}$, where each L_i is length-uniform and $\ell' \in \mathbb{N}$. Then $||L_1 \circ \dots \circ L_{m-1}||$ is a multiple of ℓ.*

The next result, again following from the "false paths" in the construction, is that in $\mathcal{T}(\langle \mathbb{A}_i^{\ell,*} \circ \mathbb{B}_s^{\ell,*} \rangle)$, Spoiler indeed needs to play the intended strategy, namely, following exactly the path a_s^ℓ in the model $\mathbb{B}_s^{\ell,*}$. This follows from the above, since there is only one path avoiding w_{trap} in model $\mathbb{B}_s^{\ell,*}$, namely the path a_s^ℓ. Hence Spoiler has to play a sequence of languages covering a_s^ℓ for each $s \in \{1,2\}^i$.

Lemma 5.4. *Let ψ be a minimal $\mathsf{ML}^{A(I)^+}$-formula equivalent to $\Box^i_{A_\ell} p$. Let $labels_\Box(v) = L_1 \dots L_m$, where each L_i is length-uniform. Let $L := L_1 \circ \dots \circ L_m$, let $d := ||L||$. Then $a_s^\ell[1 \dots d] \in L$ for each $s \in \{1,2\}^i$ such that v covers s.*

The above two results can now be used to prove that we can indeed apply the pigeonhole technique (Theorem 3.3). For this we show, in the final two results of this section, that each branch of $T^\psi(\langle \mathbb{A}_i^{\ell,*} \circ \mathbb{B}_i^{\ell,*} \rangle)$ only covers a restricted number of strings $s \in \{1,2\}^i$. The reason for this is that the application of each available operator $\Box_1$, $\Box_2$ and $\Box_{A_{\ell'}}$ for $\ell' \neq \ell$ comes with the "cost" of excluding a significant set of values s that the corresponding branch of the formula covers.

The first of these two results addresses the case where a branch a the formula equivalent to $\Box^i_{A_\ell}$ uses an operator $A_{\ell'}$, where ℓ' is not a multiple of ℓ. Due to the above Lemma 5.3, such occurrances are restricted to modal depths which themselves are a multiple of ℓ. Therefore, *immediately after* such an application, no operator $\Box_{A_{\ell''}}$ can appear, and a classical operator $\Box_j$ for $j \in \{1,2\}$ must be used. Hence such a branch can only "cover" paths in the model that have the symbol j at the next position, which is only true for half of the words in $(A_\ell)^i$. Hence each such application of a modal operator in a branch halves the number of strings s covered with this branch. One can also derive the expressiveness part of Theorem 5.1 from this lemma, since in the logic $\mathsf{ML}^{A(I)}$, the required classical operators simply are not available.

Lemma 5.5. *Let ψ be a minimal $\mathsf{ML}^{A(I)^+}$-formula equivalent to $\Box^i_{A_\ell} p$. Let $v \in T^\psi(\langle \mathbb{A}_i^{\ell,*} \circ \mathbb{B}_i^{\ell,*} \rangle)$ such that v covers s_1 and s_2. Let $labels_\Box(v) = L_1 \dots L_m$, and let $L_i = A_{f \cdot \ell + q}$ for $f \geq 0$ and $1 \leq q < \ell$. Then $s_1[u] = s_2[u]$, where $u = \frac{1}{\ell} \cdot ||L_1 \circ \dots \circ L_{i-1}|| + f + 1$.*

Our last result in this section addresses the case that an operator $A_{\ell'}$ with some "large" ℓ' (i.e., larger than ℓ) appears. In this case, the operator $A_{\ell'}$ only addresses worlds that are reachable on a path a_s^ℓ that has a sequence of ℓ' consecutive alternations. This directly implies restrictions on the string s as follows:

Lemma 5.6. *Let ψ be a minimal $\mathsf{ML}^{A(I)^+}$-formula equivalent to $\Box^i_{A_\ell} p$. Let $v \in T^\psi(\langle \mathbb{A}_i^{\ell,*} \circ \mathbb{B}_i^{\ell,*} \rangle)$ such that v covers $s \in \{1,2\}^i$, let $labels_\Box(v) = L_1 \dots L_m$, and let $L_i = A_{\ell'}$. Then, for $u = \frac{1}{\ell} ||L_1 \dots L_{i-1}||$ and all j with $1 \leq j < \frac{\ell'}{\ell}$, we have:*

- *If ℓ is even, then $s[u+j] = s[u+j+1]$.*
- *If ℓ is odd, then $s[u+j] = 3 - s[u+j+1]$.*

5.4 Applications for Expressiveness

We now obtain the expressiveness part of Theorem 2.3: If $\ell \notin I$, then the logic $\mathsf{ML}^{A(I)}$ cannot express the formula $\Box_{A_\ell} p$. (Recall that, unlike $\mathsf{ML}^{A(I)^+}$, the logic $\mathsf{ML}^{A(I)}$ does not contain the standard modal operators $\Box_1$ and $\Box_2$.) As discussed earlier, the result follows from Lemma 5.3 with a syntactic argument.

Theorem 5.7. *Lei $I \subseteq \mathbb{N}$, let $\ell \in \mathbb{N}$ with $\ell \notin I$. Then there is no formula $\varphi \in \mathsf{ML}^{A(I)}$ that is equivalent to $\Box_{A_\ell} p$.*

5.5 Applications for Succinctness

We now show that if $\ell \notin I$, then in every mimimal $\mathsf{ML}^{A(I)^+}$-formula ψ equivalent to $\Box^i_{A_\ell} p$ for an even number i, each branch can cover at most $2^{\frac{i}{2}}$ models. This upper bound is tight, since the formula $\Box^{\frac{i}{2}}_{A_{2\ell}} p$ covers $2^{\frac{i}{2}}$ strings, namely each $s \in \{11, 22\}^{\frac{i}{2}}$.

Lemma 5.8. *Let ψ be a minimal $\mathsf{ML}^{A(I)^+}$-formula equivalent to $\Box^i_{A_\ell} p$, where i is even and $\ell \notin I$. Let v be a leaf of $T^\psi(\langle \mathbb{A}_i^{\ell,*} \circ \mathbb{B}_i^{\ell,*} \rangle)$, and let $S \subseteq \{1,2\}^i$ be the set of strings s such that v covers s. Then $|S| \leq 2^{\frac{i}{2}}$.*

From Lemma 5.8 and the pigeonhole technique, we directly obtain obtain the following result.

Theorem 5.9. *Let $I \subseteq \mathbb{N}$ with $\ell \notin I$, and let $\psi \in \mathsf{ML}^{A(I)}$ be equivalent to $\Box^i_{A_\ell} p$, where i is an even number. Then $|\psi| \geq 2^{\frac{i}{2}}$.*

6 Conclusion

We proved that the expressiveness- and succinctness relationships between modal logics can be as complex as any finite or countable partial order. In the first setting we studied logics of the form $\mathsf{ML}^{\mathcal{F}}$ for a set $\mathcal{F}$ of Boolean functions. Here we obtained a complete characterization of the relative expressiveness and succinctness of logics $\mathsf{ML}^{\mathcal{F}}$ and $\mathsf{ML}^{\mathcal{G}}$. It is an interesting open question to obtain a similar complete characterization for the second setting, i.e., to answer completely the question for which sets of languages $\mathcal{L}$ and $\mathcal{K}$ we have that $\mathsf{ML}^{\mathcal{L}}$ is more succinct or more expressive than $\mathsf{ML}^{\mathcal{K}}$.

References

ABvdT10. Guillaume Aucher, Guido Boella, and Leendert van der Torre. Privacy policies with modal logic: The dynamic turn. In Guido Governatori and Giovanni Sartor, editors, *DEON*, volume 6181 of *Lecture Notes in Computer Science*, pages 196–213. Springer, 2010.

AI03. Micah Adler and Neil Immerman. An *n!* lower bound on formula size. *ACM Trans. Comput. Log.*, 4(3):296–314, 2003.

BdRV01. Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Scie*. Cambridge University Press, Cambridge, 2001.

FHMV95. Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning about Knowledge*. 1995.

FvdHIK13. Tim French, Wiebe van der Hoek, Petar Iliev, and Barteld P. Kooi. On the succinctness of some modal logics. *Artif. Intell.*, 197:56–85, 2013.

GKPS95. Goran Gogic, Henry A. Kautz, Christos H. Papadimitriou, and Bart Selman. The comparative linguistics of knowledge representation. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence, IJCAI 95, Montréal Québec, Canada, August 20-25 1995, 2 Volumes*, pages 862–869. Morgan Kaufmann, 1995.

GPT87. George Gargov, Solomon Passy, and Tinko Tinchev. Modal environment for boolean speculations. In *Mathematical logic and its applications*, pages 253–263. Springer, 1987.

HM92. J. Halpern and Y. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54(2):319–379, 1992.

vdHI14. Wiebe van der Hoek and Petar Iliev. On the relative succinctness of modal logics with union, intersection and quantification. In *AAMAS*, AAMAS '14, pages 341–348, Richland, SC, 2014. International Foundation for Autonomous Agents and Multiagent Systems.

Wil99. Thomas Wilke. CTL^{+} is Exponentially more Succinct than CTL. In C. Pandu Rangan, Venkatesh Raman, and Ramaswamy Ramanujam, editors, *FSTTCS*, volume 1738 of *Lecture Notes in Computer Science*, pages 110–121. Springer, 1999.

A Omitted Results and Definition

A.1 Facts about Formula Size Games

In [FvdHIK13], it was shown that if v is a node of a closed game tree labelled $\langle \mathbb{C} \circ \mathbb{D} \rangle$, then $\mathbb{C}$ and $\mathbb{D}$ do not contain bisimilar pointed models. In particular, this implies the following:

Proposition A.1. *Let v be a node in a closed game tree T labelled $\langle \mathbb{C} \circ \mathbb{D} \rangle$. Then $\mathbb{C} \cap \mathbb{D} = \emptyset$.*

The definition of formula-size games immediately leads to the following easy property:

Lemma A.2. *Let $T \in \mathcal{T}\langle \mathbb{A} \circ \mathbb{B} \rangle$, and let $v \in T$ be a node labelled with $\langle \mathbb{C} \circ \mathbb{D} \rangle$, and let $labels_\Box(v) = O_1 \dots O_m$. Let $(M, w) \in \mathbb{C}$, and let $(M', w') \in \mathbb{D}$. Then there exist worlds $w_0, w_1, \dots, w_m = w \in M$ and $w'_0, w'_1, \dots, w'_m = w' \in M'$ such that for each $j \in \{1, \dots, m\}$, we have that $w_j \in O_j(M, w_{j-1})$ and $w'_j \in O_j(M', w'_{j-1})$ and*

1. *(M, w_0) is an element of the class corresponding to $\mathbb{A}$,*
2. *(M', w'_0) is an element of the class corresponding to $\mathbb{B}$.*

Proof. Let r be the root of T. We show the claim by induction on the length of the path from r to v. In the base case, we have that $r = v$ and therefore $\langle \mathbb{A} \circ \mathbb{B} \rangle = \langle \mathbb{C} \circ \mathbb{D} \rangle$, and the path contains an even number of negations. Hence the claim follows trivially.

Now assume that the claim is true for the unique predecessor nove v_0 of v in T, where v_0 is labelled with $\langle \mathbb{C}_0 \circ \mathbb{D}_0 \rangle$.

We make a case distinction depending on the label of v_0. Note that since v_0 is not a leaf, v_0 cannot be labelled with an atomic proposition.

- If v_0 is labelled with $\neg$, then in particular, $labels_\Box(v_0) = labels_\Box(v)$, $\mathbb{C}_0 = \mathbb{D}$, and $\mathbb{D}_0 = \mathbb{C}$. The claim follows trivially by induction.
- If v_0 is labelled with $\vee$, then $\mathbb{C} \subseteq \mathbb{C}_0$ and $\mathbb{D} = \mathbb{D}_0$.
- Finally, let v_0 be labelled with $\Box_{O_m}$. It then follows that $labels_\Box(v_0) = O_1 \dots O_{m-1}$.
  By definition of the game, we know that for each $(M, w) \in \mathbb{C}$, there is a world $w_{m-1} \in M$ with $(M, w_{m-1}) \in \mathbb{C}_0$ and $w \in O_m(M, w_{m-1})$, and for each $(M', w') \in \mathbb{D}$, there is a world $w'_{m-1} \in M'$ with $(M', w'_{m-1}) \in \mathbb{D}_0$ and $w' \in O_m(M', w'_{m-1})$. Due to induction, worlds $w_0, w_1, \dots, w_{m-2} \in M$ and $w'_0, w'_1, \dots, w'_{m-2} \in M'$ can be chosen with the required properties. This concludes the proof.

A.2 Bisimulations

Definition A.3. *Let $M_1 = (W^1, R_1^1, \ldots, R_n^1, \Pi^1)$ and $M_2 = (W^2, R_1^2, \ldots, R_n^2, \Pi^2)$ be Kripke models, let $\mathcal{O}$ be a set of successor selection functions. A relation $Z \subseteq W^1 \times W^2$ is an* $\mathcal{O}$-bisimulation between M_1 and M_2 *if for all $(w_1, w_2) \in Z$, the following holds:*

- *for all $p \in P$, we have that $w_1 \in \Pi^1(p)$ if and only if $w_2 \in \Pi^2(p)$,*
- *(forward condition) for all $O \in \mathcal{O}$ and all $w_1' \in O(M_1, w_1)$, there is some $w_2' \in O(M_2, w_2)$ such that $(w_1', w_2') \in Z$,*
- *(back condition) for all $O \in \mathcal{O}$ and all $w_2' \in O(M_2, w_2)$, there is some $w_1' \in O(M_1, w_1)$ such that $(w_1', w_2') \in Z$.*

The following is easy to see:

Proposition A.4. *Let M_1 and M_2 be Kripke models, let $\mathcal{O}$ be a set of successor selection functions, let Z be a $\mathcal{O}$-bisimulation between M_1 and M_2, and let $(w^1, w^2) \in Z$. Then for each formula φ of $\mathsf{ML}^{\mathcal{O}}$, we have that $M_1, w^1 \models \varphi$ if and only if $M_2, w^2 \models \varphi$.*

Proof. As usual by induction on the formula. The base case where φ is a propositional variable is trivial, the cases where φ is a disjunction or a negation follow by induction. Hence let $\varphi = \Box_O \psi$, and let $M_1, w_1 \models \varphi$. To show that $M_2, w_2 \models \varphi$, let $w_2' \in O(M_2, w_2)$. Since $(w_1, w_2) \in Z$ and Z is a $\mathcal{O}$-bisimulation, there is some world $w_1' \in O(M_1, w_1)$ with $(w_1', w_2') \in Z$. Since $M_1, w^1 \models \varphi$, it follows that $M_1, w_1' \models \psi$, and hence due to induction we have that $M_2, w_2' \models \psi$. Therefore, it follows that $M_2, w_2' \models \varphi$. The converse is symmetric.

B Proofs of Results in Main Paper

B.1 Extensions of Formula Size Games

B.1.1 Proof of Theorem 3.2 The proof of Theorem 3.2 is an adaptation of the corresponding result in [FvdHIK13], the extension to arbitrary modal operators is straight-forward.

Theorem 3.2. [FvdHIK13] *Spoiler wins the FSG($\mathcal{O}$) starting with $\langle \mathbb{A} \circ \mathbb{B} \rangle$ in k moves if and only if there is a formula $\varphi \in \mathsf{ML}^{\mathcal{O}}$ with $|\varphi| = k$ such that $\mathbb{A} \models \varphi$ and $\mathbb{B} \models \neg\varphi$.*

Proof. First assume that there is a formula φ of size k such that $\mathbb{A} \models \varphi$ and $\mathbb{B} \models \neg\varphi$. We prove by induction on the construction of φ that Spoiler can win the FSG starting with $\langle \mathbb{A} \circ \mathbb{B} \rangle$ in k moves by using the strategy encoded in the formula φ.

If φ is a propositional variable p, then clearly Spoiler can win by playing the move p.

If $\varphi = \neg\psi$ for a modal formula ψ, then Spoiler plays the not-move, which results in a node labelled $\langle \mathbb{B} \circ \mathbb{A} \rangle$. Since $\mathbb{A} \models \varphi$ and $\mathbb{B} \models \neg\varphi$, it follows that

$\mathbb{A} \models \neg\psi$ and $\mathbb{B} \models \psi$. Hence due to induction, Spoiler can win the game with starting node $\langle \mathbb{B} \circ \mathbb{A} \rangle$ for the formula ψ with $|\psi|$ nodes, and thus wins the game for the formula φ with $|\varphi| = |\psi| + 1$ nodes.

If $\varphi = \psi \vee \chi$, then Spoiler chooses sets $\mathbb{A}_1$ and $\mathbb{A}_2$ with $\mathbb{A}_1 \cup \mathbb{A}_2 = \mathbb{A}$ and $\mathbb{A}_1 \models \psi$ and $\mathbb{A}_2 \models \chi$. Clearly, $\mathbb{B} \models \neg\psi$ and $\mathbb{B} \models \neg\chi$. Therefore, by induction Spoiler can win the game for ψ on $\langle \mathbb{A}_1 \circ \mathbb{B} \rangle$ in $|\psi|$ moves, and can win the game for χ on $\langle \mathbb{A}_2 \circ \mathbb{B} \rangle$ in $|\chi|$ moves. Therefore, Spoiler can win the game for φ on $\langle \mathbb{A} \circ \mathbb{B} \rangle$ with $|\psi| + |\chi| + 1 = |\varphi|$ nodes as required.

If $\varphi = \Box_O \psi$, then spoiler plays an O move as follows: Since $\mathbb{B} \models \neg\Box_O\psi$, Spoiler can choose a set $\mathbb{B}_1$ such that for each $(M, w) \in \mathbb{B}$ there is some (M, w') with $w' \in O(M, w)$ such that $M, w' \models \neg\psi$; it then follows that $\mathbb{B}_1 \models \neg\psi$. On the other hand, since $\mathbb{A} \models \Box_O\psi$, for the set $\mathbb{A}_1 = \{(M, w') \mid (M, w) \in \mathbb{A}, w' \in O(M, w)\}$, we have that $\mathbb{A}_1 \models \psi$. By induction, we therefore know that Spoiler can win the game on the mode $\langle \mathbb{A}_1 \circ \mathbb{B}_1 \rangle$ in $|\psi|$ moves, and hence can win the game on $\langle \mathbb{A} \circ \mathbb{B} \rangle$ for φ in $|\psi| + 1 = |\varphi|$ moves as required.

For the converse, assume that Spoiler can win the FSG($\mathcal{O}$) starting with node $\langle \mathbb{A} \circ \mathbb{B} \rangle$ in k moves; let T be a corresponding tree with size k. Clearly, when we only consider the labels p, $\neg$, $\vee$ and O_i, the tree T represents a formula φ from $\mathsf{ML}^{\mathcal{O}}$ with $|\varphi| = k$. By induction, we prove that for each node v labelled with $\langle \mathbb{A} \circ \mathbb{B} \rangle$ in T, for the formula φ_v represented by the subtree corresponding to v, we have that $\mathbb{A} \models \varphi_v$ and $\mathbb{B} \models \neg\varphi_v$.

If v is a leaf, then v is labelled with a propositional variable p. Due to the winning condition, we know that v is closed, hence $\mathbb{A} \models p$ and $\mathbb{B} \models \neg p$. Now assume that v is not a leaf, then v is labelled with $\neg$, $\vee$, or some O_i for $i \in \{1, \dots, k\}$.

First assume that v is labelled with $\neg$. Then v has a single successor node u labelled with $\langle \mathbb{B} \circ \mathbb{A} \rangle$, and $\varphi_v = \neg\varphi_u$. By induction, we know that $\mathbb{B} \models \varphi_u$, and $\mathbb{A} \models \neg\varphi_u$. Hence $\mathbb{A} \models \varphi_v$ and $\mathbb{B} \models \neg\varphi_v$ as required.

Now assume that v is labelled with $\vee$, then v has two successor nodes u_1 and u_2 with $\varphi_v = \varphi_{u_1} \vee \varphi_{u_2}$ labelled with $\langle \mathbb{A}_1 \circ \mathbb{B} \rangle$ and $\langle \mathbb{A}_2 \circ \mathbb{B} \rangle$ with $\mathbb{A}_1 \cup \mathbb{A}_2 = \mathbb{A}$. By induction, we know that $\mathbb{A}_1 \models \varphi_{u_1}$, $\mathbb{A}_2 \models \varphi_{u_2}$, $\mathbb{B} \models \neg\varphi_{u_1}$ and $\mathbb{B} \models \neg\varphi_{u_2}$. Therefore, each pointed model $(M, w) \in \mathbb{A}$ satisfies φ_{u_1} or φ_{u_2}, it follows that $\mathbb{A} \models \varphi_{u_1} \vee \varphi_{u_2} = \varphi$, and each pointed model $(M, w) \in \mathbb{B}$ satisfies $\neg\varphi_{u_1}$ and $\neg\varphi_{u_2}$, hence $\mathbb{B} \models \neg(\varphi_{u_1} \vee \varphi_{u_2}) = \varphi$ as required.

Finally assume that v is labelled with O. Then v has a unique successor u, and $\varphi_v = O\varphi_u$ for some $i \in \{1, \dots, k\}$, and u is labelled with $\langle \mathbb{A}_1 \circ \mathbb{B}_1 \rangle$, where $\mathbb{A}_1 = \{(M, w') \mid (M, w) \in \mathbb{A}, w' \in O_i(M, w)\}$, and for each $(M, w) \in \mathbb{B}$, there is a pointed model $(M, w') \in \mathbb{B}_1$ with $w' \in O(M, w)$. Due to induction, we know that $\mathbb{A}_1 \models \varphi_u$, and $\mathbb{B}_1 \models \neg\varphi_u$. By the choice of $\mathbb{A}_1$ and $\mathbb{B}_1$, it therefore follows that $\mathbb{A} \models O\varphi_u = \varphi_v$, and $\mathbb{B} \models \neg O\varphi_u = \neg\varphi_v$ as required.

B.1.2 Proof of Theorem 3.3 (pigeonhole principle)

Theorem 3.3. *Let $\mathcal{O}$ be a set of successor selection functions. Let φ be a formula, let $\mathbb{A}$ and $\mathbb{B}$ be sets of pointed models such that $\mathbb{A} \models \varphi$ and $\mathbb{B} \models \neg\varphi$. If for every nontrivial leaf u of every closed game tree $T_{\mathcal{O}} \in \mathcal{T}(\langle \mathbb{A} \circ \mathbb{B} \rangle)$, the class*

of models corresponding to $\mathbb{A}$ ($\mathbb{B}$) has size at most c, then every $\mathsf{ML}^{\mathcal{O}}$-formula equivalent to φ has size at least $\frac{|\mathbb{A}|}{c}$ ($\frac{|\mathbb{B}|}{c}$).

Proof. Assume that this is not the case, and let there be a formula ψ that is equivalent to φ. In particular, then $\mathbb{A} \models \psi$ and $\mathbb{B} \models \neg\psi$. Hence let $T^{\psi}_{(}\langle \mathbb{A} \circ \mathbb{B}\rangle)$ be the closed game tree $T \in \mathcal{T}(\langle \mathbb{A} \circ \mathbb{B}\rangle)$ that corresponds to playing the strategy ψ. Clearly, $T^{\psi}_{(}\langle \mathbb{A} \circ \mathbb{B}\rangle)$ is isomorphis to ψ.

Then, by the proof of Theorem 3.2, each node of ψ corresponds to a lead in $T^{\psi}_{(}\langle \mathbb{A} \circ \mathbb{B}\rangle)$, and each pointed model from $\mathbb{A}$ ($\mathbb{B}$) appeads in at least one leaf of $T^{\psi}_{(}\langle \mathbb{A} \circ \mathbb{B}\rangle)$. Since there are only c pointed models in each leaf of $T^{\psi}_{(}\langle \mathbb{A} \circ \mathbb{B}\rangle)$, the tree $T^{\psi}_{(}\langle \mathbb{A} \circ \mathbb{B}\rangle)$ has at least $\frac{|\mathbb{A}|}{c}$ ($\frac{|\mathbb{B}|}{c}$). Therefore, φ has at least this many leafs as well, and in particular, the size of φ is at least $\frac{|\mathbb{A}|}{c}$ ($\frac{|\mathbb{B}|}{c}$) as claimed.

B.2 Proofs of Results in Section 4 (Single Step Operators)

B.2.1 Proof of Theorem 4.1

Theorem 4.1. *Let $\mathcal{F}$ and $\mathcal{G}$ be sets of n-ary Boolean functions. Then the following are equivalent:*

1. $\mathsf{ML}^{\mathcal{G}} \leq_{expr} \mathsf{ML}^{\mathcal{F}}$,
2. *for each $g \in \mathcal{G}$, there is a set $S \subseteq \mathcal{F}$ such that $g \equiv \bigvee_{f\in S} f$.*

Proof. Since there are only a finite number of Boolean functions of each arity, let $\mathcal{F} = \{f_1, \ldots, f_k\}$, and let $\mathcal{G} = \{g_1, \ldots, g_l\}$. The direction 2 to 1 is trivial: If $g \equiv f_{i_1} \vee \cdots \vee f_{i_t}$, then $M, w \models \Box_g\varphi$ if and only if $M, w' \models \varphi$ for all worlds w' such that w' is an I-successor of w for some I with $f_{i_m}(I) = 1$ for one of the i_m. Therefore, $\Box_g\varphi$ is equivalent to $\Box_{f_{i_1}}\varphi \wedge \cdots \wedge \Box_{f_{i_t}}\varphi$.

It remains to show that if one of the g_is is not of this form, then there is a formula φ of $\mathsf{ML}^{\mathcal{G}}$ that cannot be expressed in $\mathsf{ML}^{\mathcal{F}}$. A standard technique to prove such results are bisimulations, which we adapt to logics of this form (see details in Appendix A.2).

Hence assume indirectly that $\Box_g p$ can be expressed in $\mathsf{ML}^{\mathcal{F}}$ and g is not of the form $\bigvee_{i\in I} f_i$ for any set $I \subseteq \{1, \ldots, k\}$. Let $S \subseteq \{1, \ldots, k\}$ be the (possibly empty) set of indices i such that f_i implies g (i.e., if $f_i(I) = 1$, then $g(I) = 1$). By choice of S, it follows that for each $i \notin S$, there is some assignment I_i with $f_i(I_i) = 1$ and $g(I_i) = 0$. It also follows that $\bigvee_{i\in S} f_i$ implies g. Since we assumed that g is not of the form $\bigvee_{i\in S} f_i$ for any S, it then follows that g does not imply $\bigvee_{i\in S} f_i$. Therefore, there is an assignment I_g such that $g(I_g) = 1$, and for each $i \in S$, we have that $f_i(I_g) = 0$.

Now consider the following models M_1 on the left-hand side and M_2 on the right-hand side:

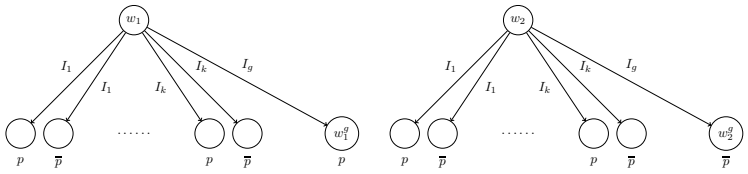

Here, an arrow labelled I_i between worlds w and w' represents that $(w, w') \in R_j$ for exactly those j with $I_i(r_j) = 1$. It is obvious that $M_1, w_1 \models \Box_g p$: Since $g(I_i) = 0$ for all $i \in \{1, \ldots, k\}$, the world w_1^g is the only g-successor of w_1 in M_1, and by definition, $M_1, w_1^g \models p$. On the other hand, $M_2, w_2 \models \neg\Box_g p$, since the only g-successor of w_2 in M_2 is w_2^g, and by definition, $M_2, w_2^g \models \overline{p}$.

We define the relation Z as follows: Z contains the pair (w_1, w_2) and all pairs of unnamed worlds in which p has the same value. We show that Z is an $\{f_1, \ldots, f_k\}$-bisimulation. For the forward condition, let w_1' be an f_i-successor of w_1. We distinguish two cases:

- if w_1' is *not* the world w_1^g, then we can simply choose w_2' to be the corresponding world w_2' (i.e., the one in the same position in the picture) of model M_2, which is then an f_i-successor of w_2 with $(w_1', w_2') \in Z$.
- if w_1' is the world w_1^g, then in particular, w_1^g is an f_i-successor of w_1. It follows that $f_i(I_g) = 1$, and therefore, $i \notin S$. By the choice of I_i, it follows that $f(I_i) = 1$. Therefore, we can choose w_2' as the I_i-successor w_2 in M_2 where p is false.

The backward condition is shown analogously.

Now indirectly assume that there is a formula φ of $\mathsf{ML}^{\mathcal{F}}$ which is equivalent to $\Box_g p$. Then in particular it follows that $M_1, w_1 \models \varphi$ and $M_2, w_2 \not\models \varphi$. However, since the above-constructed bisimulation Z contains the pair (w_1, w_2) and φ is a formula from $\mathsf{ML}^{\{f_1,\ldots,f_n\}}$, it follows from Proposition A.4 that $M_1, w_1 \models \varphi$ if and only if $M_2, w_2 \models \varphi$. Hence we have a contradiction.

B.2.2 Proof of Theorem 4.2

Theorem 4.2. *Let $\mathcal{F}$ and $\mathcal{G}$ be sets of Boolean functions such that $\mathsf{ML}^{\mathcal{F}}$ and $\mathsf{ML}^{\mathcal{G}}$ are equally expressive and $\mathcal{G} \not\subseteq \mathcal{F}$. Then $\mathsf{ML}^{\mathcal{G}}$ is exponentially more succinct than $\mathcal{F}$.*

Proof. We prove the slightly stronger result that as soon as $\mathcal{G}$ contains a function g which is a disjunction of functions in $\mathcal{F}$, but not an element of $\mathcal{F}$, then the formula $\neg\Box_g^i\neg p$ needs exponential length when expressed as an $\mathsf{ML}^{\mathcal{F}}$-formula (an equivalent $\mathsf{ML}^{\mathcal{F}}$-formula does exist due to Theorem 4.1). For n-ary Boolean functions f_1 and f_2, we write $f_1 \leq f_2$ if $f_1(r_1, \ldots, r_n) \leq f_2(r_1, \ldots, r_n)$ for all $r_1, \ldots, r_n \in \{0, 1\}$. We define

- $\mathcal{F}_1 = \{f \in \mathcal{F} \mid f \leq g\}$, and
- $\mathcal{F}_2 = \mathcal{F} \setminus \mathcal{F}_1$.

Since g is a disjunction of functions in $\mathcal{F}$, it then clearly follows that $g = \vee_{f \in \mathcal{F}_1} f$. Let $\mathcal{F}_1 = \{f_1, \ldots, f_k\}$, and let $\mathcal{F}_2 = f_{k+1}, \ldots, f_m$. For each $i \in \{k+1, \ldots, m\}$, let $\overrightarrow{\alpha_i} = (r_1^i, \ldots, r_n^i)$ be chosen such that $f_i(\overrightarrow{\alpha_i}) = 1$, and $g(\overrightarrow{\alpha_i}) = 0$. Such a sequence exists since $f_i \not\leq g$ for $i \in \{k+1, \ldots, m\}$.

Further, let a set of vectors $\left\{\overrightarrow{\beta_1}, \ldots \overrightarrow{\beta_t}\right\}$ be a smallest set chosen such that

- $g(\overrightarrow{\beta_1}) = \cdots = g(\overrightarrow{\beta_t}) = 1$,
- there is no $f \in \mathcal{F}_1$ with $f(\overrightarrow{\beta_1}) = \cdots = f(\overrightarrow{\beta_t}) = 1$.

Such a set exists, since for each $f \in \mathcal{F}_1$ we have that $f \leq g$ and $f \neq g$ (since $g \notin \mathcal{F}$). In particular, choosing $\overrightarrow{\beta_1}, \ldots \overrightarrow{\beta_t}$ as the set of all assignments $\overrightarrow{\beta}$ with $g(\overrightarrow{\beta}) = 1$ satisfies the two conditions (although not minimality). Clearly, for the smallest such set, we still have that $t \geq 2$, since for each $\overrightarrow{\beta}$ with $g(\overrightarrow{\beta}) = 1$ there is some $f \in \mathcal{F}_1$ with $f(\overrightarrow{\beta}) = 1$, since we know that $g = \vee_{f \in \mathcal{F}_1} f$.

For each $i \in \mathbb{N}$, let $\varphi_i = \neg \Box_g^i \neg p$. Clearly, φ_i is a $\mathsf{ML}^{\mathcal{G}}$-formula (in fact even a $\mathsf{ML}^{\{g\}}$-formula) and the length of φ_i is linear in i. Since g is a disjunction of functions in $\mathcal{F}$, due to Theorem 4.1, it follows that for each i, there is some $\mathsf{ML}^{\mathcal{F}}$-formula ψ_i of minimal length such that φ_i and ψ_i are equivalent. To prove the theorem, it suffices to show that the length of each ψ_i is at least $(\frac{t}{t-1})^i$.

To show this, we construct models similarly to the ones from the proof in [FvdHIK13]. Our models are based on trees of width t and depth i, and are constructed as follows:

- Each tree T has a root w_0 with depth 0.
- Each node $u \in T$ with depth smaller than i has successors $v_1^u, \ldots, v_t^u$, where v_i^u is a $\overrightarrow{\beta_i}$-successor of u. (Note that due to the minimality of $\overrightarrow{\beta_1}, \ldots \overrightarrow{\beta_t}$, the sequence consists of pairwise different vectors).

In this proof only, for a word $s = s_1 \ldots s_i \in \{1, \ldots, t\}^*$, we say that a node $u \in T$ is an s-successor of a world v if $s = \epsilon$ and $u = v$, or if there is an intermediate node u' such that u' is a $\overrightarrow{\beta_{s_1}}$-successor of v and u is (inductively) a $s_2 \ldots s_i$-successor of u'.

We now define our models as follows:

- For each $s \in \{1, \ldots, t\}^i$, let A_s be the model obtained from the tree T, where in the unique world w_s that is an s-successor of the root of T, the variable p is true.
- The model B is the model obtained from the tree T, where the variable p is false in every world.

Additionally, if u and v are nodes with $depth(v) = depth(u) + 1$ and one of the following is true:

- u is a node of some A_s and v is a node of B, or
- v is a node of some A_s and u is a node of B,

then v is an α_j-successor of u for each $j \in \{k+1, \ldots, m\}$.

Let $\mathbb{A} = \left\{ A_s \mid s \in \{1, \ldots, t\}^i \right\}$, and $\mathbb{B} = \{B\}$. Then:

- $\mathbb{A} \models \varphi_i$, since the world w_s satisfies the variable p, and $g(\overrightarrow{\beta_j}) = 1$ for all relevant j,

– $\mathbb{B} \models \neg\varphi_i$, since the model B does not contain any world in which p is true and which can be reached on a path adressed by $\Box^i_g$.

We first show that no $T \in \mathcal{T}(\langle \mathbb{A} \circ \mathbb{B} \rangle)$ can contain a nontrivial node (i.e., a node labelled $\langle \mathbb{C} \circ \mathbb{D} \rangle$ with $\emptyset \notin \{\mathbb{C}, \mathbb{D}\}$) that is labelled with $\Box_{f_j}$ for $j \in \{k+1, \dots, m\}$ (i.e., $f_j \in \mathcal{F}_2$). Recall from above that in this case, $g(\overrightarrow{\alpha_j}) = 0$, and $f_j(\overrightarrow{\alpha_j}) = 1$. Assume indirectly that such a node u labelled with $\langle \mathbb{C} \circ \mathbb{D} \rangle$ exists. Let the successor node of u be labelled with $\langle \mathbb{C}_1 \circ \mathbb{D}_1 \rangle$.

We make a case distinction:

– First assume that $\mathbb{C}$ corresponds to $\mathbb{A}$ and $\mathbb{D}$ corresponds to $\mathbb{B}$. Then $\mathbb{C}_1$ contains, in particular, all $\overrightarrow{\alpha_j}$-successors of all nodes in $\mathbb{C}$, which includes all nodes in B of the corresponding depth. In particular, this includes the successor picked for the right-hand side in the model B. Therefore, we have a contradiction to Proposition A.1.
– The second case is symmetric.

We theorefore know that the formulas ψ_i do not contain any occurrance of an operator $\Box_f$ for $f \in \mathcal{F}_2$, hence ψ_i is in fact a $\mathsf{ML}^{\mathcal{F}_1}$-formula. Hence to conclude the proof, it suffices to show that every leaf in a tree $T \in \mathcal{T}(\langle \mathbb{A} \circ \mathbb{B} \rangle)$ contains at most $(t-1)^i$ elements, the result then follows from Theorem 3.3, since $|\mathbb{A}| = t^i$. To show this, let $T \in \mathcal{T}(\langle \mathbb{A} \circ \mathbb{B} \rangle)$, and let u be a leaf in T. Then $labels_{\Box}(u) = f_{j_1} \dots f_{j_i}$ (clearly, the modal depth of ψ_i must be i), with $j_1, \dots, j_i \in \{1, \dots, k\}$. By construction for each f_{j_l}, there is one value $\overrightarrow{\beta_h}$ among $\overrightarrow{\beta_1}, \dots, \overrightarrow{\beta_t}$ with $f_{j_l}(\overrightarrow{\beta_h}) = 0$. Therefore, for each of the t successors of each node in each A_s, the application of f_{j_l} covers at most $t-1$ many of them. Since the depth of the formula (and the tree T) is i, this implies that each leaf contains only at most $(t-1)^i$ many of the models A_s.

B.2.3 Proof of Main Result of Single-Step Operators, Theorem 2.2

We now use the characterizations of expressiveness and succinctness obtained in Theorems 4.1 and 4.2 to prove our main result on single-step operators, Theorem 2.2.

Theorem 2.2. *Let S be a finite set, and let $\leq_S$ be a partial order on S. Then there exist families of sets of $\lceil \log_2(|S|+1) \rceil$-ary Boolean functions $(\mathcal{F}_s)_{s \in S}$ and $(\mathcal{G}_s)_{s \in S}$ such that for each $s, t \in S$, the following holds:*

1. $\mathsf{ML}^{\mathcal{F}_s} \leq_{expr} \mathsf{ML}^{\mathcal{F}_t}$ *if and only if* $\mathsf{ML}^{\mathcal{G}_s} \leq_{poly} \mathsf{ML}^{\mathcal{G}_t}$ *if and only if* $s \leq_S t$.
2. *All logics* $\mathsf{ML}^{\mathcal{G}_s}$ *are equally expressive, and if* $s \not\leq_S t$, *then* $\mathsf{ML}^{\mathcal{G}_s}$ *is exponentially more succinct that* $\mathsf{ML}^{\mathcal{G}_t}$.

Proof. Let $S = \{s_1, \dots, s_n\}$, with an ordering chosen such that if $s_i \leq_S s_j$, then $i \leq j$. Let $k = \lceil \log_2(|S|+1) \rceil$, then $2^k > |S|$. Hence there is an injective function $i \colon S \to \mathcal{P}(\{1, \dots, k\})$, such that $i(s) \neq \emptyset$ for all $s \in S$.

We first prove the succinctness result, i.e., define the sets $\mathcal{F}_s$ for $s \in S$. For this, we use the k modalities $\Box_1, \dots, \Box_k$. Let $\mathcal{P}$ contain all projections, i.e., all

k-ary Boolean functions of the form $p_i(r_1, \ldots, r_k) = r_i$ for some $i \in \{1, \ldots, k\}$. We now define, for each $s \in S$, the function $f_s = \vee_{j \in i(s)} r_j$, and then define $\mathcal{F}_{s_i}$ inductively (recall that if $s_j \leq_S s_i$, then $j \leq i$) as

$$\mathcal{F}_{s_i} = \mathcal{P} \cup \{f_{s_i}\} \cup \bigcup_{s_j \leq_S s_i} \mathcal{F}_{s_i}.$$

Since all involved Boolean functions are disjunctions of functions in $\mathcal{P}$, and each $\mathcal{F}_{s_j}$ contains $\mathcal{P}$ as a subset, it follows from Theorem 4.1 that all $\mathsf{ML}^{\mathcal{F}_s}$ are equally expressive as $\mathsf{ML}^{\mathcal{P}}$. In particular, all $\mathsf{ML}^{\mathcal{F}_s}$ are equally expressive. By construction, if $s \leq_S t$, then $\mathcal{F}_s \subseteq \mathcal{F}_t$, and hence in particular, every $\mathsf{ML}^{\mathcal{F}_s}$-formula is also a $\mathsf{ML}^{\mathcal{F}_t}$-formula as claimed.

Now assume that $s \not\leq_S t$, and let $s = s_i$, $t = s_j$ for some $i, j \in \{1, \ldots, n\}$. By construction, it follows that the function f_s is an element of $\mathcal{F}_{s_i}$, but not an element of $\mathcal{S}_{s_j}$. Since $\mathsf{ML}^{\mathcal{F}_{s_i}}$ and $\mathsf{ML}^{\mathcal{F}_{s_j}}$ are equally expressive, Theorem 4.2 then implies that $\mathsf{ML}^{\mathcal{F}_{s_i}}$ is exponentially more succinct than $\mathsf{ML}^{\mathcal{F}_{s_j}}$. This completes the proof.

For the expressiveness result, we use a very similar construction, but leave out the projections (as their role was to ensure that all logics have the same expressive power). We define the function $g_s = \oplus_{j \in i(s)} r_j$, and define the sets $\mathcal{G}_s$ as follows (inductively as above):

$$\mathcal{G}_{s_i} = \{g_s\} \cup \bigcup_{s_j \leq_S s_i} \mathcal{G}_{s_i}.$$

The proof is identical to the succinctness case above, since Theorem 4.1 implies that $\Box_{f_s}$ cannot be expressed with any number of opeators $\Box_{f_{s'}}$ for $s' \neq s$.

B.3 Proofs of Results in Section 5 (Arbitrary-Step Operators)

B.3.1 Modal Depth For a node $v \in T$ where $labels_{\Box}(v) = L_1 \ldots L_m$ for length-uniform languages $L_1, \ldots, L_m$, we say that the *modal depth* of v is the value $||L_1 \circ \cdots \circ L_m||$. We denote this value with $md(v)$. A straight-forward induction on the path from the root to the node v shows the following:

Proposition B.1. *Let $v \in T \in \mathcal{T}_{\mathcal{O}}(\mathbb{A}_i^{\ell,*}, \mathbb{B}_i^{\ell,*})$, where $\mathcal{O}$ contains only length-uniform languages. Let $(\mathbb{X}, w_X)$ be covered by v, where $\mathbf{X} \in \{\mathbf{A}, \mathbf{B}\}$. Then $depth(w_X) = md(v)$.*

We say that a model M is *complete*, if every world $w \in M$ has both a 1- and a 2-successor. Note that all models $\mathbb{A}_i^{\ell,*}$ and $\mathbb{B}_s^{\ell,*}$ are complete, but $\mathbb{A}_i^{\ell}$ and $\mathbb{B}_s^{\ell}$ are not. We say that a node v *covers* a pair of models (M_A, M_B) if v is labelled $\langle \mathbb{C} \circ \mathbb{D} \rangle$ and there are worlds $w_A \in M_A$ and $w_B \in M_B$ such that $(M_A, w_A) \in \mathbb{C}$ and $(M_B, w_B) \in \mathbb{D}$ or $(M_A, w_A) \in \mathbb{D}$ and $(M_B, w_B) \in \mathbb{C}$. In particular, then v is a subformula of φ that distinguishes (M_A, w_A) and (M_B, w_B) in the sense that $M_A, w_a \models \varphi$ if and only if $M_B, w_b \models \neg\varphi$.

B.3.2 Proof of Lemma 5.2

Lemma 5.2. *Let $I \subseteq \mathbb{N}$, let ψ be a minimal $\mathsf{ML}^{A(I)^+}$-formula equivalent to $\Box^i_{A_\ell} p$, let v be a node of $T^\psi(\langle \mathbb{A}_i^{\ell,*} \circ \mathbb{B}_i^{\ell,*} \rangle)$ labelled $\langle \mathbb{A} \circ \mathbb{B} \rangle$, and let $(M, w) \in \mathbb{A} \cup \mathbb{B}$. Then $w \neq w_{trap}$.*

Proof. Since ψ is equivalent to $\Box^i_{A_\ell} p$, we know that ψ is p-monotone in the following sense: If M and M' are models where M' is obtained from M by making p true in additional worlds and $M, w \models \psi$, then $M', w \models \psi$ holds as well. Since ψ is minimal, this implies that ψ only contains positive occurrences of p, i.e., the variable p only occurs under an even numer of negations. Furthermode, p is the only propositional variable appearing in ψ. Therefore, every leaf of $T^\psi(\langle \mathbb{A}_i^{\ell,*} \circ \mathbb{B}_i^{\ell,*} \rangle)$ is labelled with the variable p and a class of models $\langle \mathbb{C} \circ \mathbb{D} \rangle$ such that $\mathbb{C}$ corresponds to $\mathbb{A}_i^{\ell,*}$, and $\mathbb{D}$ to $\mathbb{B}_i^{\ell,*}$.

We now show inductively that for every node $v \in T^\psi(\langle \mathbb{A}_i^{\ell,*} \circ \mathbb{B}_i^{\ell,*} \rangle)$ labelled with $\langle \mathbb{C} \circ \mathbb{D} \rangle$ or $\langle \mathbb{D} \circ \mathbb{C} \rangle$ such that $\mathbb{C}$ corresponds to $\mathbb{A}_i^{\ell,*}$ and $\mathbb{D}$ corresponds to $\mathbb{B}_i^{\ell,*}$, for every model $(\mathbb{A}_i^{\ell,*}, w_A) \in \mathbb{C}$ there is a descendent of w_A in $\mathbb{A}_i^{\ell,*}$ where p is true, and for every model $(\mathbb{B}_s^{\ell,*}, w_B) \in \mathbb{D}$, there is a descendent of w_B in $\mathbb{B}_s^{\ell,*}$ where p is false. Since the reflexive singleton of $\mathbb{A}_i^{\ell,*}$ ($\mathbb{B}_s^{\ell,*}$) does not have a successor where p is true (false), this shows that $T^\psi(\langle \mathbb{A}_i^{\ell,*} \circ \mathbb{B}_i^{\ell,*} \rangle)$ does not contain a node that is labelled with a model $(\mathbb{A}_i^{\ell,*}, w_{\text{trap}})$ or $(\mathbb{B}_s^{\ell,*}, w_{\text{trap}})$.

We prove this claim inductively over the tree structure. For the leaves, the claim follows from the above, as every leaf is labelled with p and $\langle \mathbb{C} \circ \mathbb{D} \rangle$ where $\mathbb{C}$ ($\mathbb{D}$) corresponds to $\mathbb{A}_i^{\ell,*}$ ($\mathbb{B}_i^{\ell,*}$), hence every pointed models in $\mathbb{C}$ satisfy p, and all pointed models in $\mathbb{B}$ satisfy $\overline{p}$.

Now let v be a non-leaf node labelled $\langle \mathbb{C} \circ \mathbb{D} \rangle$ in $T^\psi(\langle \mathbb{A}_i^{\ell,*} \circ \mathbb{B}_i^{\ell,*} \rangle)$ such that v is not a leaf and the above claim is true for all successors of v. We make a case distinction:

- If v is labelelled $\neg$, then v has a single successor v' for which the claim holds by induction. The result for v follows trivially, since v' is labelled $\langle \mathbb{D} \circ \mathbb{C} \rangle$ or $\langle \mathbb{C} \circ \mathbb{D} \rangle$, where $\mathbb{D}$ corresponds to $\mathbb{B}_i^{\ell,*}$ and $\mathbb{C}$ to $\mathbb{A}_i^{\ell,*}$.
- If v is labelled $\vee$, then v has two successors labelled $\langle \mathbb{C}_1 \circ \mathbb{D} \rangle$ and $\langle \mathbb{C}_2 \circ \mathbb{D} \rangle$ with $\mathbb{C}_1 \cup \mathbb{C}_2 = \mathbb{D}$, and for which the claim is true. Since the claim is true for each model in $\mathbb{C}_1$ and $\mathbb{C}_2$, it is also true for their union, $\mathbb{C}$.
- If v is laballed $\Box_L$, then v has a single successor v' labelled $\langle \mathbb{C}_1 \circ \mathbb{D}_1 \rangle$, where $\mathbb{C}_1$ ($\mathbb{D}_1$) contains (at least) one descendent for each pointed model in $\mathbb{C}$ ($\mathbb{D}$). Since the claim is true for the sets $\mathbb{C}_1$ and $\mathbb{D}_1$ and the descendent relation is transitive, the claim for $\mathbb{C}$ and $\mathbb{D}$ follows.

This completes the proof.

B.3.3 Proof of Lemma 5.3

Lemma 5.3. *Let ψ be a minimal $\mathsf{ML}^{A(I)^+}$-formula equivalent to $\Box^i_{A_\ell} p$. Let $v \in T^\psi(\langle \mathbb{A}_i^{\ell,*} \circ \mathbb{B}_i^{\ell,*} \rangle)$. Let $labels_\Box(v) = L_1 \dots L_{m-1} A_{\ell'}$, where each L_i is length-uniform and $\ell' \in \mathbb{N}$. Then $||L_1 \circ \dots \circ L_{m-1}||$ is a multiple of ℓ.*

Proof. Assume that this is not the case, and let $d := ||L_1 \circ \cdots \circ L_{m-1}||$, let v be labelled $\langle \mathbb{C} \circ \mathbb{D} \rangle$. By Proposition B.1, for every pointed model $(M, w) \in \mathbb{C} \cup \mathbb{D}$, we have that $depth(w) = d$. By construction, if d is not a multiple of ℓ, then each world $w \in \mathbb{A}_i^{\ell,*}$ or $\mathbb{B}_s^{\ell,*}$ with $depth(w) = d$ does not have both a 1- and a 2-successor in the base model $\mathbb{A}_i^{\ell}$ ($\mathbb{B}_s^{\ell}$). Since $A_{\ell'}$ contains the word $a_1^{\ell'}$ (starting with 1) and $a_2^{\ell'}$ (starting with 2), the successor node v' of v, labelled with $\langle \mathbb{C}_1 \circ \mathbb{D}_1 \rangle$, contains a model $(\mathbb{A}_i^{\ell,*}, w_A)$or $(\mathbb{B}_s^{\ell,*}, w_B)$ in $\mathbb{C}_1$ where $w_A \notin \mathbb{A}_i^{\ell}$ or $w_B \notin \mathbb{B}_s^{\ell}$, i.e., $w_A = w_{\text{trap}}$ or $w_B = w_{\text{trap}}$. This is a contradiction to Lemma 5.2.

B.3.4 Proof of Lemma 5.4

Lemma 5.4. *Let ψ be a minimal $\mathsf{ML}^{A(I)^+}$-formula equivalent to $\Box^i_{A_\ell} p$. Let $labels_\Box(v) = L_1 \dots L_m$, where each L_i is length-uniform. Let $L := L_1 \circ \cdots \circ L_m$, let $d := ||L||$. Then $a_s^\ell[1 \dots d] \in L$ for each $s \in \{1,2\}^i$ such that v covers s.*

Proof. Due to Lemma 5.2, we know that for every covered model $(\mathbb{B}_s^{\ell,*}, w_B)$, w_B is a world of the base model $\mathbb{B}_s^{\ell}$. Due to Proposition B.1, we know that for $(\mathbb{B}_s^{\ell,*}, w_B)$ as above, $depth(w_B) = d$. Therefore, w_B is the unique $a_s^\ell[1 \dots d]$-successor of $\mathbb{B}_s^{\ell,*}$'s root in $\mathbb{B}_s^{\ell,*}$. Hence, due to Lemma A.2, we know that $a_s^\ell[1 \dots d] \in L$.

B.3.5 Proof of Lemma 5.5

Lemma 5.5. *Let ψ be a minimal $\mathsf{ML}^{A(I)^+}$-formula equivalent to $\Box^i_{A_\ell} p$. Let $v \in T^\psi(\langle \mathbb{A}_i^{\ell,*} \circ \mathbb{B}_i^{\ell,*} \rangle)$ such that v covers s_1 and s_2. Let $labels_\Box(v) = L_1 \dots L_m$, and let $L_i = A_{f \cdot \ell + q}$ for $f \geq 0$ and $1 \leq q < \ell$. Then $s_1[u] = s_2[u]$, where $u = \frac{1}{\ell} \cdot ||L_1 \circ \cdots \circ L_{i-1}|| + f + 1$.*

Proof. Since $L_i = A_{f \cdot \ell + q}$ with $q \geq 1$, Lemma 5.3 implies that $||L_1 \circ \cdots \circ L_{i-1}|| = g \cdot \ell$ for some g. Let $u = \frac{1}{\ell} \cdot ||L_1 \circ \cdots \circ L_{i-1}|| + f + 1 = g + f + 1$ and let $d' := ||L_1 \circ \cdots \circ L_i||$, then $d' = g \cdot \ell + f \cdot \ell + q' = (g + f) \cdot \ell + q'$. Since this is not a multiple of ℓ, Lemma 5.3, implies that L_{i+1} (and $i + 1$ does exist, since otherwise the formula does not have the full modal depth) cannot be $A_{\ell'}$ for any ℓ', hence $L_{i+1} = \{\alpha\}$ for some $\alpha \in \{1, 2\}$. Then it follows that for each word $x \in L$, we have that $x[d' + 1] = x[(g + f) \cdot \ell + q' + 1] = \alpha$. Due to Lemma 5.4, we know that $a_{s_1}^\ell, a_{s_2}^\ell \in L$. With the above, this implies that $a_{s_1}^\ell[(g + f) \cdot \ell + q' + 1] = a_{s_2}^\ell[(g + f) \cdot \ell + q' + 1] = \alpha$.

Now indirectly assume that $s_1[u] \neq s_2[u]$, i.e., $s_1[g + f + 1] \neq s_2[g + f + 1]$. In particular, then $a_{s_1}^\ell[(g + f) \cdot \ell + q' + 1] \neq a_{s_2}^\ell[(g + f) \cdot \ell + q' + 1]$, which is a contradiction to the above.

B.3.6 Proof of Lemma 5.6

Lemma 5.6. *Let ψ be a minimal $\mathsf{ML}^{A(I)^+}$-formula equivalent to $\Box^i_{A_\ell} p$. Let $v \in T^\psi(\langle \mathbb{A}_i^{\ell,*} \circ \mathbb{B}_i^{\ell,*} \rangle)$ such that v covers $s \in \{1,2\}^i$, let $labels_\Box(v) = L_1 \dots L_m$, and let $L_i = A_{\ell'}$. Then, for $u = \frac{1}{\ell} ||L_1 \dots L_{i-1}||$ and all j with $1 \leq j < \frac{\ell'}{\ell}$, we have:*

- *If ℓ is even, then $s[u+j] = s[u+j+1]$.*
- *If ℓ is odd, then $s[u+j] = 3 - s[u+j+1]$.*

Proof. Withous loss of generality, we assume $\ell\text{“} > \ell$, since otherwise, there is no j in the required interval and the claim is trivial. With d_i, we again denote $||L_1 \circ \cdots \circ L_i||$. Since $L_i = A_{\ell'}$ with $\ell' \geq 1$, we know from Lemma 5.3 that d_{i-1} is a multiple of ℓ, and by choice of u it follows that $d_{i-1} = u \cdot \ell$. In particular, u is a natural number.

Let $L = L_1 \circ \cdots \circ L_m$. Then, due to Lemma 5.4, we know that $a^\ell_s[1 \dots d_m] \in L$.

Since $L_i = A_{\ell'}$, we know that for each word $x \in L$, the subword $x[d_{i-1} + 1 \dots d_{i-1} + \ell']$ is alternating.

Therefore, since $a^\ell_s \in L$ and $d_{i-1} = u \cdot \ell$, we know that $a^\ell_s[u \cdot \ell + 1 \dots u \cdot \ell + \ell']$ is alternating, i.e., for each position $i \in \{u \cdot \ell + 1, \dots, u \cdot \ell + \ell' - 1\}$, we have that $a^\ell_s[i] \neq a^\ell_s[i+1]$.

Now let $1 \leq j < \frac{\ell'}{\ell} 1$, and let $i = (u+j) \cdot \ell$. Then, since $j \geq 1$ and $\ell \geq 1$, it follows that $i = (u+j) \cdot \ell \geq (u+1) \cdot \ell \geq u \cdot \ell + 1$, and since $j < \frac{\ell'}{\ell}$, we have that $i = (u+j) \cdot \ell < (u + \frac{\ell'}{\ell}) \cdot \ell = u \cdot \ell + \ell'$, and hence $i \leq u \cdot \ell + \ell' - 1$. Therefore, i is in the above interval, and hence $a^\ell_s[i] \neq a^\ell_s[i+1]$.

Since $i = (u+j) \cdot \ell$, by the definition of a^ℓ_s, it follows that $a^\ell_s[i]$ is the last symbol of $a^\ell_{s[u+j]}$, and $a^\ell_{s[i+1]}$ is the first symbol of $a^\ell_{s[u+j+1]}$.

Hence we know that the last symbol of $a^\ell_{s[u+j]}$ is different from the first symbol of $a^\ell_{s[u+j+1]}$.

- If ℓ is even, then for both $\alpha \in \{1, 2\}$, the first symbol of a^ℓ_α is α, and the last symbol of a^ℓ_α is $3 - \alpha$. Hence $\alpha = s[u+j] \neq s[u+j+1] = 3 - \alpha$ would imply that the last symbol of $a^\ell_{s[u+j]}$ (namely $3 - \alpha$) is identical to the first symbol of $a^\ell_{s[u+j+1]}$ (which is also $3 - \alpha$), but from the above we know that the last symbol of $a^\ell_{s[u+j]}$ is different from the first symbol of $a^\ell_{s[u+j+1]}$. Hence in this case $s[u+j] = s[u+j+1]$.
- If ℓ is odd, then for both $\alpha \in \{1, 2\}$, both the first and the last symbol of a^ℓ_α is α. Since the last symbol of $a^\ell_{s[u+j]}$ is different from the first symbol of $a^\ell_{s[u+j+1]}$, this implies that $s[u+j] \neq s[u+j+1]$. Since only the symbols 1 and 2 appear, this means that $s[u+j] = 3 - s[u+j+1]$.

B.3.7 Proof of Theorem 5.7

Theorem 5.7. *Lei $I \subseteq \mathbb{N}$, let $\ell \in \mathbb{N}$ with $\ell \notin I$. Then there is no formula $\varphi \in \mathsf{ML}^{A(I)}$ that is equivalent to $\Box_{A_\ell} p$.*

Proof. Assume that such a formula exists, and let ψ be one of minimal size. Define classes of models $\mathbb{A} = \mathbb{A}^{\ell,*}_1$ and $\mathbb{B} = \mathbb{B}^{\ell,*}_1$. Then $\mathbb{A} \models \psi$ and $\mathbb{B} \models \neg\psi$. Let $T = T^\psi(\langle \mathbb{A}^{\ell,*}_i \circ \mathbb{B}^{\ell,*}_i \rangle)$. Then T only contains the operators available in $\mathsf{ML}^{A(I)}$. Clearly, there is a leaf v of T that covers the string $s = 1$. Clearly, every ancestor of v covers the string s as well. Let $labels_\Box(v) = L_1 \dots L_m$. Then for each i,

we have that L_i is of the form $A_{\ell'}$ for some $\ell' \neq l$. In particular, we have that $L_1 = A_{\ell'}$ for some $\ell' \neq \ell$. Clearly, we can without loss of generality assume that $\ell' < \ell$. From Lemma 5.3, it then follows that L_2 cannot be of the form $A_{\ell'}$ for any $\ell' \in \mathbb{N}$. However, since in the logic $\mathsf{ML}^{A(I)}$ with $\ell \notin I$, only languages of this form occur, we have a contradiction.

B.3.8 Proof of Lemma 5.8

Lemma 5.8. *Let ψ be a minimal $\mathsf{ML}^{A(I)^+}$-formula equivalent to $\Box^i_{A_\ell} p$, where i is even and $\ell \notin I$. Let v be a leaf of $T^\psi(\langle \mathbb{A}_i^{\ell,*} \circ \mathbb{B}_i^{\ell,*} \rangle)$, and let $S \subseteq \{1,2\}^i$ be the set of strings s such that v covers s. Then $|S| \leq 2^{\frac{i}{2}}$.*

Proof. Let $labels_\Box(v) = L_1 \dots L_m$. We consider each j for which $||L_1 \circ \dots \circ L_{j-1}||$ is a multiple of ℓ, and show how the operator L_i restricts the possible values of s. Hence let $||L_1 \circ \dots \circ L_{j-1}|| = g \cdot \ell$, we say that $g \cdot \ell$ is the depth in which this operator appears. There are three cases to consider.

1. If $L_i = \{\alpha\}$ for some $\alpha \in \{1,2\}$, then L contains only words x with $x[g \cdot \ell + 1] = \alpha$. Due to Lemma 5.4, we know that $a^\ell_s \in L$ for all $s \in S$, and hence $a^\ell_s[g \cdot \ell + 1] = \alpha$, which implies that $s[g+1] = \alpha$. Hence this operator rules out $\frac{1}{2}$ of all possible strings in $\{1,2\}^i$. Due to Lemma 5.3, the languages $L_{i+1}, \dots, L_{i+\ell-2}$ are not of the form $A_{\ell'}$ for some ℓ', hence the next restriction occurs at depth $(g+1) \cdot \ell$.
2. If $L_i = A_{f \cdot \ell}$ for some $f > 1$ (recall that $L_i \neq A_\ell$ for all i), then, by Lemma 5.6, the elements of S must satisfy a sequence of $(f-1)$ equalities. Hence this operator rules out all but $\frac{1}{2^{(f-1)}}$ strings in $\{1,2\}^i$, and, again due to Lemma 5.3, the next restriction appears at depth $(g+f) \cdot \ell$.
3. If $L_i = A_{f \cdot \ell + q}$ for some $f \geq 0$ and $1 \leq q < \ell$, then, by Lemma 5.6, the elements of S must satisfy a sequence of f identities (one identity for each $j \in \{0, \dots, f-1\}$). Hence all but $\frac{1}{2^f}$ elements of $\{1,2\}^i$ are ruled out, and, as above, the next restriction appears at the next multiple of ℓ, i.e., at depth $(g+f+1) \cdot \ell$.

Note that is is easy to see that the conditions required by L_i at different indices are independent, as they refer to different indices of the strings s. Hence the following three operations appear:

- Increase depth by ℓ, and add a restriction factor of $\frac{1}{2}$,
- Increase depth by $f \cdot \ell$, and add a restriction factor of $\frac{1}{2^{f-1}}$,
- Increase depth by $(f+1) \cdot \ell$, and add a restriction factor of $\frac{1}{2^f}$.

In each case, increasing the depth by $2 \cdot \ell$ adds a restrictin factor of at least $\frac{1}{2}$. Since the complete depth must be $i \cdot \ell$, this means that the minimum restriction factor is at most $\frac{1}{2^{\frac{i}{2}}}$, i.e., we have that $|S| \leq 2^i \cdot \frac{1}{2^{\frac{i}{2}}} = 2^{\frac{i}{2}}$ as claimed.

B.3.9 Proof of Theorem 5.9

Theorem 5.9. *Let $I \subseteq \mathbb{N}$ with $\ell \notin I$, and let $\psi \in \mathsf{ML}^{A(I)}$ be equivalent to $\Box^i_{A_\ell} p$, where i is an even number. Then $|\psi| \geq 2^{\frac{i}{2}}$.*

Proof. Recall that $\mathbb{B}_i^{\ell,*} = \left\{\mathbb{B}_s^{\ell,*} \mid s \in \{1,2\}^i\right\}$. In particular $|\mathbb{B}| = 2^i$, and that $\mathbb{A}_i^{\ell,*} \models \psi$ and $\mathbb{B}_i^{\ell,*} \models \neg\psi$. Due to Theorem 3.3, it thus suffices to show that each leaf u of a tree $T \in \mathcal{T}(\langle \mathbb{A}_i^{\ell,*} \circ \mathbb{B}_i^{\ell,*} \rangle)$ covers at most $2^{\frac{i}{2}}$ strings $s \in \{1,2\}^i$. Clearly it is enough to show the result for $T = T^{\psi}(\langle \mathbb{A}_i^{\ell,*} \circ \mathbb{B}_i^{\ell,*} \rangle)$, since ψ is a formula equivalent to $\Box^i_{A_\ell}$ of minimal size. Without loss of generality we can assume that the modal depth of each leaf ov T is exactly $i \cdot \ell$. Clearly, for each $s \in \{1,2\}^i$, there is a leaf v_s of T that covers s. Lemma 5.8 states that each leaf v_s can cover at most $2^{\frac{i}{2}}$ elements as claimed. Therefore, T must have at least $\frac{2^i}{2^{\frac{i}{2}}} = 2^{\frac{i}{2}}$ leaves, which concludes the proof.

B.3.10 Proof of Main Result on Alternation Languages, Theorem 5.1

Theorem 5.1. *Let $I_1, I_2 \subseteq \mathbb{N}$ with $I_1 \nsubseteq I_2$. Then*

1. *$\mathsf{ML}^{A(I_1)^+}$ is exponentially more succinct than $\mathsf{ML}^{A(I_2)^+}$, and*
2. *there is an $\mathsf{ML}^{A(I_1)}$-formula for which there is no equivalent $\mathsf{ML}^{A(I_2)}$-formula.*

Proof. Since $I_1 \nsubseteq I_2$, there is some $\ell \in I_1 \setminus I_2$.

1. For each $i \in \mathbb{N}$, define $\varphi_i = \Box^{2i}_{A_\ell} p$. Then clearly, the length of φ_i is linear in i. For each i, let ψ_i be the smallest formula in $\mathsf{ML}^{A(I_2)}$ that is equivalent to φ_i. Then, due to Theorem 5.9, we know that $|\psi_i| \geq 2 \cdot 2^{\frac{i}{2}}$, hence the length of ψ_i is exponential in the length of φ_i as claimed.
2. The formula $\Box_{A_\ell} p$ is a $\mathsf{ML}^{A(I_1)}$-formula, and due to Theorem 5.7, there is no $\mathsf{ML}^{A(I_2)}$-formula equivalent to $\Box_{A_\ell} p$.

B.3.11 Proof of Main Result on Arbitrary-step Operators, Theorem 2.3

Theorem 2.3. *Let S be a countable set, and let $\leq_S$ be a partial order on S. Then there exist families of languages $(\mathcal{L}_s)_{s\in S}$ and $(\mathcal{K}_s)_{s\in S}$ over the alphabet $\{1,2\}$ such that for each $s,t \in S$, the following holds:*

1. *$\mathsf{ML}^{\mathcal{K}_s} \leq_{expr} \mathsf{ML}^{\mathcal{K}_t}$ if and only if $\mathsf{ML}^{\mathcal{L}_s} \leq_{poly} \mathsf{ML}^{\mathcal{L}_t}$ if and only if $s \leq_S t$.*
2. *All logics $\mathsf{ML}^{\mathcal{L}_s}$ are equally expressive, and if $s \not\leq_S t$, then $\mathsf{ML}^{\mathcal{L}_s}$ is exponentially more succinct that $\mathsf{ML}^{\mathcal{L}_t}$.*

Proof. The proof is very similar to the proof of Theorem 2.2. Let $S = (s_\ell)_{\ell\in\mathbb{N}}$, where $s_j \leq_S s_i$ implies $j \leq i$. Now, for $\ell \geq 1$, inductively define I_{s_ℓ} as follows:

$$I_{s_\ell} = \{\ell\} \cup \bigcup_{t \leq_S s} I_t.$$

Clearly, we have that $I_s \subseteq I_t$ if and only if $s \leq_S t$. Therefore, the result follows from Theorem 5.1 with the choice $\mathcal{L}_s = \{A_\ell \mid \ell \in I_s\} \cup \{\{1\}, \{2\}\}$ and $\mathcal{K}_s = \{A_\ell \mid \ell \in I_s\}$.